Data Base Management Systems. 1976

A Special Datapro Report

datapro

About Datapro Research Corporation

Datapro Research Corporation is the most widely accepted and respected source of up-to-date, cost-saving information about data processing and office products and services. The company was founded in 1968 to do high technology research and consulting. In January 1970, **Datapro 70** was delivered to charter subscribers. Since then, the **Datapro 70** service has come to be regarded as "the EDP buyer's bible" in well over 10,000 subscriber sites around the world. The company subsequently compiled and released **Datapro Reports On Minicomputers**, **Datapro Reports On Banking Automation**, **Datapro Reports On Office Systems**, **Datapro Directory Of Software**, **Datapro Reports On Retail Automation**, **Datapro Reports On Data Communications**, and **Datapro EDP Solutions**. Although the Datapro reference services are designed primarily to aid users, they also have widespread acceptance among equipment manufacturers, software companies, consultants, financial analysts, and educators.

Now in its ninth successful year, Datapro Research Corporation serves over 30,000 subscribers, delivering up-to-date, comprehensive information about data processing, data communications, and office systems.

For those who value information.

datapro

Datapro Research Corporation, 1805 Underwood Boulevard, Delran, NJ 08075 609/764/0100 A McGraw-Hill Company

Phoenix AZ (602) 263-7831 Washington DC (301) 933-4020 Atlanta GA (404) 636-8021 Chicago IL (312) 644-2000

Recd 26 My 77

Data Base Management Systems. 1976

A Special Datapro Report

This report is one of several hundred such reports on data processing and office system hardware, software, services and companies that make up the authoritative Datapro volumes. These volumes are an integral part of each of Datapro's four-part information services for EDP and office professionals. The other service components, subscribed to on an annual basis, include monthly supplements to the volumes, monthly interpretive newsletters, and Custom Consulting with our analysts. Completely independent in its research and evaluations, Datapro publishes the most widely used EDP reference and information services.

datapro

DATAPRO RESEARCH CORPORATION 1805 Underwood Boulevard, Delran, New Jersey 08075, (609) 764-0100

NOTICES

When US Government drawings, specifications, or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use or sell any patented invention that may in any way be related thereto.

Qualified requesters may obtain copies from DDC, Document Service Center, Cameron Station, Alexandria, Virginia, 22314. Orders will be expedited if placed through the Librarian or other person designated to request documents from DDC.

Do not return this copy unless return is required by security considerations, contractual obligations, or notice on a specific document.

A Buyer's Guide to Data Base Management Systems

This report presents useful guidelines to aid you in selecting and implementing a data base management system, plus comparison charts to help you assess the key characteristics of 13 leading systems: ADABAS, DATACOM/DB, DL/1 DOS/VS, DMS-II, DMS/90, DMS-1100, IDMS, I-D-S/II, IMS, INQUIRE, MODEL 204, SYSTEM 2000, and TOTAL.

According to all the EDP journals and publications, the future is now insofar as data base management systems are concerned. There is more current emphasis on this type of software than on any other, with the possible exception of data communications monitors. In fact, users' interests in these two areas are closely knit. Nearly every vendor of a leading data base management system also offers a complementary data communications monitor.

Most of the systems described in this report were originally designed to operate on an IBM System/360 or 370 computer, but many of these can also be used on one or more of the other major mainframes. In addition, this year we have broadened this report to include the data base management systems developed by Burroughs, Honeywell, and Sperry Univac for operation on their respective mainframes. Most of the systems can be installed to run under almost any operating system offered with a given manufacturer's computers. The DBMS is typically designed to be independent of the operating system—but, in fact, most systems are sensitive to logic changes in any operating system version.

Since the advent of data base concepts, DBMS software has been the object of a great many questions and a great deal of controversy. We at Datapro are acutely aware of this inasmuch as the subject regularly generates more inquiries to our Telephone Consulting Service than any other software topic.

Since you are reading this report, it's likely that you are either interested in moving to a data base management environment or in reviewing what you already have installed in your company. If you are thinking of installing a DBMS, you'll need some guidelines to assist you in your planning effort, and then in evaluating the many packages now on the market. In this report, we'll try to outline the steps required to embark on this endeavor with a minimum of frustration and cost. This report is not designed to be a tutorial on theoretical concepts, but rather to give some sound advice based on the experience of our staff and the many users we've interviewed who have already implemented a DBMS.

As mentioned previously, this report deals with 13 of the most prominent data base management systems being marketed at this time. We would like to remind the reader that there are other packages on the market—including DBMS packages designed specifically for minicomputers—which may be justifiable within your installation, and we certainly don't want to exclude them arbitrarily. The guidelines presented here are general guidelines that can be used to evaluate and plan for any DBMS. The effort involved in installing a minicomputer DBMS can be as large as for a full-blown DBMS on a giant mainframe; the degree of preparation is not necessarily proportionate to the size of the mainframe. The extent of implementation of applications programs may or may not be smaller, but the planning efforts are much the same.

DBMS or DMS?

Before going any further in this report, it's necessary to clarify what is meant by *data base management systems* as opposed to a similarly named but distinctly different class of software often called *data management systems*. These two terms have been thrown around rather loosely within the industry, and have led to massive confusion on the part of the uninitiated. Many are the times that a Datapro subscriber has called our Telephone Consulting Service to ask advice about a DBMS, only to realize, after talking to one of our consultants, that what he really wants to install is a data management system.

A *data base management system* can be defined as a software system that is intended to manage and maintain data in a non-redundant structure for the purpose of being processed by *multiple applications*. A data base management system organizes data elements in some predefined structure, and retains relationships between different data elements within the data base.

A *data management system*, on the other hand, is one that is intended primarily to permit access to, and retrieval from, already existing files (usually for a *single application*). Although a data management system may provide the capability to minimize data redundancy and centralize the storage of data, the principal intent of the system is to perform such functions as information retrieval, report generation, and inquiry for a single application. This category of products includes such packages as Informatics, Inc.'s MARK IV (Report 70E-500-01) and Mathematica's RAMIS (Report 70E-610-01), which have all the qualifications of true data base management systems (they do, in fact, have their own file structures) except for the fact that their primary intended use is for the processing of single application files.

Some products that started out as data management systems have grown with the times to the point where they have become full-blown data base management systems. An example of this is a product from Infodata Systems called INQUIRE. Through a series of enhancements and restructurings, this system has grown into a bona-fide DBMS.

ABSTRACT

This report describes an application of Computer-Aided Design concepts to the general preliminary design of ships, in which shape description plays an important part. Although the present study is preliminary in nature and will require considerable elaboration for practical use, it does indicate the feasibility of the approach.

The design and evaluation of hull forms was accomplished "on line" using the Project MAC time-shared digital computer, the display console developed by the MIT Electronic Systems Laboratory, and a very general, parametric surface description technique developed by Professor S.A. Coons. Three-dimensional hull surfaces displayed on the CRT screen could be altered in a few seconds by typed-in changes in parameters, and could be rotated to any desired viewing angle for study. Using these techniques, the lines of the US DD 692 were simulated such that the routines for calculating midships coefficient, prismatic coefficient, displacement, wetted surface area, and centers of buoyancy yielded values closely resembling those of the actual ship. A brief economic analysis shows the great saving in time and cost and the corresponding increase in study of alternative designs which would be possible with such a design system.

A Buyer's Guide to Data Base Management Systems

In this report we will deal with the more widely known and used data base management systems: ADABAS, DATACOM/DB, DL/1 DOS/VS, DMS-II, DMS/90, DMS-1100, IDMS, I-D-S/II, IMS, INQUIRE, MODEL 204, SYSTEM 2000, and TOTAL. We acknowledge that there may be many other products that are every bit as good and may match our definition of a DBMS, and we don't intend for you to overlook these systems in your evaluation. In addition, the vendors of the above-named products are continually announcing new subsets of their full-blown systems for use with smaller equipment configurations. These subsets carry titles such as "entry level" or carry suffix numbers to identify the subsets. The purpose of these subsets is clearly to introduce a form of DBMS to smaller users who cannot afford all the frills and options of the larger systems which have gained the most attention and popularity. These more limited types of systems will be further discussed later in this report.

The Make-Up of a DBMS

Prior to the introduction of data base management systems, every computer installation was forced to employ an array of programmers and systems people to define and implement their required applications. Each of these specialists had to be completely cognizant of how the various applications programs interfaced with the operating system, the data, and the computer environment within which they were used.

The DBMS concept offers an intermediate buffer between the applications programmer and these environmental considerations. With a DBMS, the applications programmer no longer has to worry about the location and structure of the data with which he must work. With most DBMS systems, the languages supported are those most commonly used (COBOL, PL/1, FORTRAN, Assembler), so the programmer can concentrate on turning out programs to access and process data without the severe constraints placed on earlier generations of programmers. A DBMS also allows non-programming personnel to utilize the data located in the data base with reduced effort. Some DBMS vendors have even created query languages designed specifically to meet the needs of non-programming users of the system. These query languages operate efficiently within the confines of the DBMS and the operating system.

During the early years of data processing, every application program utilized files that were unique to that particular application, leading to the storage and manipulation of a great deal of redundant data. Each operation within the corporate structure created its own files in the format required for its unique functional activities. The DBMS approach was introduced to alleviate these redundancies and permit greater efficiency in the storage and utilization of information. The key concept employed in DBMS systems is that data can be tied together in such a manner as to allow any application program the ability to access any or all parts of the data, regardless of location, access method, or record make-up.

During the last few years, it has been widely recognized that the concept of a generalized and centralized corporate data base is the logical approach to faster and more efficient operations.

Most DBMS offerings on the market today are classified as generalized data base management systems, although currently there are several vendors of data base management systems who are deviating from the generalized approach to a more specialized DBMS approach. This can be seen more readily in the evolution of minicomputer DBMS products, time-sharing operations, and service bureau facilities. There is little doubt that this specialized approach will gradually gain in popularity.

A generalized data base management system should incorporate most, and preferably all, of the following facilities:

- Application program independence from the DBMS control programs.
- Support of the programming language(s) used in the corporate environment prior to installation.
- Utility programs to facilitate creation and maintenance of the data base(s).
- Facilities for data reorganization.
- The ability to effect data security and access limitation.
- Automatic restart capabilities in case of system failure, or the ability to recover operations manually with minimum effort.
- System facilities for "fine tuning" of the DBMS.

Most generalized data base management systems offer their users the opportunity to make performance versus storage trade-offs. For users interested in implementing a system that is geared for fast retrieval, the DBMS will make certain storage allocations designed to minimize the I/O demands of the system, thereby necessarily increasing the main storage demands. The reverse is true for users desiring to minimize their storage requirements; in this case, more accessing time will necessarily be needed by the system. To optimize the DBMS configuration for each installation, most vendors offer a series of options designed to "fine tune" their systems. In fact, what these options offer is a means of balancing the trade-offs to achieve maximum performance in a particular environment. The fine-tuning responsibility is usually assigned to an individual referred to as the Data Base Administrator (DBA).

As mentioned before, some vendors are now offering subsets of their full-scale DBMS products. Along with the main purpose of introducing these cut-down versions for use on smaller hardware systems, some vendors have designed the subset versions to maximize the performance of the inquiry/response functions in the users' applica-

A Buyer's Guide to Data Base Management Systems

tions. The subsets are normally upward-compatible with the larger versions and require comparatively little effort to expand the operational capabilities of the corporate data base. The basic features of the full-blown versions are usually found in the smaller (or entry-level) versions, but the latter are designed to handle smaller data bases on smaller hardware configurations. The subset systems are usually somewhat, but not much, easier to install, and there are normally fewer fine-tuning options available to the Data Base Administrator.

Learn What It's All About

Before entering the first phase of DBMS selection, it's important to know exactly what a data base management system consists of and what it can do. To achieve this knowledge, if you have not had any prior exposure to the subject, we recommend that you avail yourself of one of the many generalized DBMS seminars now being offered (such as those sponsored by Datapro). We must emphasize, however, that attending a seminar conducted *by a vendor* of a DBMS system will not expose you to an impartial overview of data base management in general. The vendors' seminars may be excellent when you have reached the point of evaluating specific systems for selection, but they're the wrong place to get your basic education in the subject.

In addition to the seminar approach, there are a number of good books available on the subject, and such publications as *Datamation* and *Computerworld* frequently run excellent articles on DBMS by industry experts.

The Planning Phase

When you can confidently tell the difference between a data base management system and a data management system, when you understand the interrelationships between a DBMS and a data communications monitor, and when you have firmly convinced yourself that a DBMS environment is what you really need, you'll be ready to embark on the most critical phase of the operation—the planning phase.

The key to planning is keeping any one person or one department from being placed into the position of having to make all the key decisions. Planning for a data base management system is a combined effort requiring the participation of almost every department in the corporation that will initially or eventually partake of the benefits of the DBMS. Of course, one person has to be in a position to direct the planning effort, and in many instances this person is a senior systems analyst. Many companies now refer to this individual as a Data Base Administrator (DBA). The DBA is the one to whom the responsibility of pulling together all the strings in the implementation cycle will fall. Other important responsibilities of the DBA include research into what is needed in the DBMS, day-to-day maintenance of the DBMS, and fine tuning of the system to maximize its performance.

During the planning phase, the DBA cannot take full responsibility for defining the needs of the organization, but must work in close coordination with representatives of all the affected departments and activities. The needs of the individual users must be defined, at least in general, in order to develop a basic implementation plan.

Some advocates of DBMS planning and installation believe that it is a waste of time to try to plan a total DBMS right from scratch. In some respects they are correct, but we believe that at least a general overall plan *must* be laid out. It must also be remembered that the initial plan will undoubtedly change as the system is being implemented. The reason why an overall plan is needed is simple—cost evaluation. Not knowing the eventual scope of the project obviously precludes any accurate cost estimating. No doubt there will be cost overruns—it is rare when there aren't—but careful planning will at least minimize the overruns.

Which leads us to point out that two of the most heavily involved people on the planning committee will be the representative from the corporate finance group and the representative of the data processing department—the former to hold the reins on the dollars, and the latter, if for no other reason, because his department will absorb most of the major impacts of the initial implementation. The data processing department will, in many cases, install new equipment, initiate new operating procedures, work with new staff personnel (such as the DBA), and institute new standards for the programming staff.

As the planning phase gets into high gear, many corporations have noted that the total operations of the corporation become much clearer to the various department representatives than ever before. This results in such comments as "I didn't know that such a function was being performed by anyone else in the corporation," or "Are you interested in that type of information, too? I thought my department was the only one using that data." This type of communication is a healthy sign in the planning phase, inasmuch as it assures the DBA that there are, in fact, records and files that will be common to many departments, and that the users will be more likely to accept the changes that are bound to occur in their methods of accessing data.

When the general planning of the data bases and their contents and formats has been thoroughly researched, then the DBA, along with assistance from the data processing group, will have to determine the type of *data base structure* best suited for the DBMS operations needed by the corporation. There are no hard and fast guidelines for making this determination. There are systems which employ hierarchical structures, chained pointer techniques, networking techniques, etc. The technique to use is the one that fits your data structure requirements, satisfies the system performance demands, and minimizes the time and effort required for the installation and cutover process.

A Buyer's Guide to Data Base Management Systems

Another factor which must be considered in the planning phase is the inclusion or exclusion of a *data communications monitor*. Contrary to what some people think, the installation of a DBMS does not necessarily demand the inclusion of a data communications network, nor does the implementation of a company-wide data communications system necessitate the installation of a DBMS. If however, your plans do call for communications, the requirements of the monitor (or a home-grown system to perform similar communications control functions) must be outlined at this time. You should find Datapro's corresponding *Buyer's Guide to Data Communications Monitors* (Report 70E-010-63) helpful in this effort.

In summary, the planning phase should yield the following items:

- A thorough definition of the corporate information needs, existing applications, data file structures and contents, retrieval requirements, and data security needs.
- A general plan that outlines the overall functions and requirements of the new system.
- A list of functions that must be provided by the DBMS to address the immediate needs of the system.
- A list of functions that the DBMS must have in order to meet the long-range plans of the corporation.
- A list of additional features that would be "nice to have" in a DBMS in order to make life easier for the corporation and afford some "above and beyond" capabilities to the various user departments.
- A clear assignment of responsibilities, not only as they apply to the Data Base Administrator, but as they apply to the entire user organization and the data processing group as well.

When it comes to selecting a specific DBMS, the functions on the "nice to have" list should be the *only* functions which you'll be willing to sacrifice or compromise. Compromises are not uncommon; in fact, when the actual implementation cycle starts, there will be changes made to the original plans almost from the beginning. Planning sessions will become a way of life, not only until the system has been implemented, but from then on, to make sure that the growth of the system is in line with the real needs of the corporation. Only with management's full backing and encouragement, as well as its active participation, will the installation of a corporate data base management system be successful and effective.

The Evaluation Phase

Now, finally, you are ready to evaluate the products offered by the many vendors of DBMS packages. Keep in mind that the system you decide upon will be the one that you will probably have to live with for a long time. You'll soon be nearly as "locked in" to your DBMS supplier as you are to your computer vendor. Implementing a DBMS system can be analogous to the wild growth of ivy; unless kept under complete control, it can quickly get out of hand and overrun the entire edifice. Keep in mind, also, that before you commit yourself there is always a fallback: preserving your existing system. Consider what would be required to enhance your existing system to match your newly defined requirements, and don't eliminate this option until you are convinced that it is not a viable alternative. By all means, don't fall victim to rapidly spreading opinion that if you do not install a DBMS, "you're just not with it."

Now is the time to contact the vendors of the data base management systems addressed in this report, as well as any others that you may be aware of. The reliability of the vendor should always be of concern to you. You want to be sure that the vendor is here to stay and will dedicate his efforts to your problems and desires.

You should closely review the comparison charts on the following pages, read the applicable Datapro product reports, review the information you receive from the vendors, and compare the capabilities of each system to your lists of functional requirements (immediate, long-range, and "nice to have"). At this point you may be able to eliminate some of the candidates because they cannot satisfy a critical requirement. The cost factor may eliminate others, especially where it would be necessary for you and/or the vendor to modify or extend the standard facilities of a DBMS. Another important factor to consider is the cost of any required hardware that does not currently exist in your installation. Remember, it is rare that any existing computer system will contain all the equipment required for a particular DBMS, but you should try to minimize additional purchases.

Next comes the detailed evaluation. A suggested procedure is as follows:

- Match the facilities of each remaining candidate (as well as your existing system) against the requirements established during the planning phase.
- Compute the estimated total cost of installation of each system, and the projected cost savings after the system is installed.
- Determine the number and skill levels of the people that will be required to support the system both during implementation and thereafter. (Incidentally, experienced DBMS users agree that it's likely to be highly beneficial to hire at least one person who's had prior experience with data base management in general, and preferably with the particular DBMS you select.)
- Run benchmarks. This can be a costly process for the vendor, but a highly illuminating one for you, and well-established vendors are generally prepared to satisfy your demands.
- Get the vendor to give in-house demonstrations geared to both your DP personnel and your user personnel.

- Review the documentation of each system for clarity and understanding.
- Develop reasonable time schedules for implementation.
- Talk to other users of the system and learn their first-hand experiences with it. If you can find a user with applications similar to the ones you are planning, it will be particularly beneficial, but *any* user experience will give you valuable additional insights into the product.
- If possible, visit some existing installations and talk to the people responsible for installing, maintaining, and using the system.

Once these comparisons and schedules have been worked up, the final decision should be relatively straightforward. What's more, it should be a decision you can live with comfortably for some time to come.

Words of Caution

Before ending this report, it is only fair to issue a few words of caution:

- If you have done your homework, planned the system thoroughly, selected the appropriate DBMS, and implemented the system properly, the rewards will be great. However, do not minimize the magnitude of the commitment required, nor the costs involved.
- It cannot be overemphasized that *experienced people* are the key to successful planning and implementation. If you have experienced technicians available, make sure you use them wisely. The Data Base Administrator (or chief analyst) should be a fully experienced person. If you do not have this caliber of people involved in the effort, it would be wise to go outside your organization and hire at least one. Eagerness is no substitute for experience, and you're going to need all the experience you can get.
- In selecting a Data Base Administrator, care must be exercised. This key individual will have the ultimate responsibility for control and organization of all the corporate data included in the data base. In a small organization, this function may be a part-time job for one person. In other, larger organizations, it may be a department, with one DBA and several other people reporting to him or her. It is important, though, that this be a staff function, not reporting to the manager of data processing. Because the DBA has to interface equally with the users as well as the DP department, and respond to the requests of all involved parties, the position should be independent of reporting responsibilities to any one group.
- Implementation should be carried out in carefully planned segments. Your applications programs do not have to be implemented all at once. In fact, it's usually an unwise move to do so. Bring up a relatively easy existing application or a new one first. Then work out a priority schedule for developing and converting the rest of your applications.
- Training should not be overlooked. Implementing a DBMS requires comprehensive training of both DP and user department personnel who will be involved in the implementation effort. This training should be scheduled on an orderly basis; it's usually not feasible to retrain an entire organization at one time.
- Keep the user departments up to date on the status of the implementation effort. Remember, they have to undergo many changes, and they will be more amenable to these changes if they feel that they are part of the total effort.
- Finally, if you should determine that you have made an error in your decision to implement a DBMS, have the guts and good sense to admit it before the effort has gone too far. Don't throw good money after bad. Corporate well-being must be your major concern.

The Comparison Charts

The charts on the following pages are designed to provide you with an overview of the comparative features of the major data base management systems on the market today. The data in the charts is meant to be used in conjunction with the individual DATAPRO 70 report on each of the systems. (The first entry on the chart will point you to the appropriate report number.) The individual reports will provide you with detailed explanations of all the features of these systems, as well as the ratings assigned by their users on the basis of extensive field experience.

Most of the entries in the charts should be self-explanatory to anyone who's completed his or her basic education in DBMS concepts. A few of the entries, however, may need the brief explanations that follow:

- *Data base organization*—defines the technique used to organize the elemental structure of the data bases. In most cases, the techniques are hierarchical, network, relational, or a combination thereof.
- *System accounting facilities*—describes the method used to capture usage statistics for charge-back purposes.
- *Concurrent batch/on-line*—indicates whether the system can support simultaneous on-line and batch operations.
- *Concurrent application program access*—indicates whether the system allows more than one application program to access a data base concurrently.
- *Price*—indicates the basic system purchase price quoted by the vendor at the time this report was prepared. Most vendors also offer various lease plans and optional features. For details, please consult the individual DATAPRO 70 report on each DBMS, as referenced in the first chart entry.□

A Buyer's Guide to Data Base Management Systems

COMPARING THE LEADING DATA BASE MANAGEMENT SYSTEMS

System	ADABAS	DATACOM/DB	DL/1 DOS/VS
DATAPRO 70 report number	70E-757-01	70E-180-02	70E-491-01
Vendor	Software AG of North America	Computer Information Management Co.	IBM
Current number of users	Over 150	30	Datapro estimates approximately 150
CONFIGURATION CPU's supported	S/360, 370; Univac 70 and 9000 Series; Siemens 4004	System/360, 370	S/370 Models 125, 135, and 145
Operating systems	DOS/VS, OS, OS/VS, MVS; TDOS, DOS; PBS, BS1000	DOS, DOS/VS, OS, OS/VS	DOS/VS
Minimum memory requirements	170K bytes	DOS—31K bytes plus buffer space	98K bytes—batch; 163K bytes—with CICS
DATA BASE FEATURES Data base organization	Network with full inversion	Inverted file	Hierarchical (sequential and direct)
Applicaton languages	COBOL, PL/1, FORTRAN, Assembler	Any language with a CALL facility	COBOL, PL/1, Assembler
Data base languages	ADABAS	DATAQUERY	Data Language/1
Access methods supported	BDAM	Proprietary access method	VSAM and SAM
Variable-length segments	Yes	Yes	Yes
Data base security	Password for access and update; encryption	Password, terminal validation	Password lockout at data set level
System accounting facilities	Report Utility for space statistics	File access and buffer usage statistics	Stand-alone—No; with CICS—Yes
RECOVERY FEATURES Checkpoint/restart	Yes	Automatic at OPEN	No
Data base integrity	Up to 15 levels of data protection	Verification of DB at OPEN time	Transaction logging for backout and recovery
OTHER SYSTEM FEATURES Concurrent batch/on-line	Yes	Yes	With CICS
Concurrent application program access	Yes	Yes	With CICS
Inquiry/retrieval facility	ADASCRIPT query language	Yes	No
Report generator	ADAWRITER	Interface to standard report writers	No
Data dictionary support	Interface to DATA-MANAGER (from MSP)	Directory only, not complete dictionary	No
Telecommunication interfaces	Most popular TP monitors	DATACOM/DC and others	CICS/DOS/VS
PRICE (basic system, no options)	DOS—$80,000; OS—$120,000	DOS, DOS/VS—$26,600; OS, OS/VS—$32,000	Monthly license only: $330/month
COMMENTS	Supports up to 255 files, up to 500 fields per file. Offers over 100 diagnostics	Up to 16 data elements can be accessed in a single CALL. Runs under EDOS	DL/1 Entry superset also available. Fine-tuning utilities included

A Buyer's Guide to Data Base Management Systems

COMPARING THE LEADING DATA BASE MANAGEMENT SYSTEMS (Continued)

System	DMS-II	DMS/90	DMS 1100	I-D-S/11
DATAPRO 70 report number	70E-112-01	70E-877-01	70E-877-02	70E-480-01
Vendor	Burroughs	Sperry Univac	Sperry Univac	Honeywell
Current number of users	Approximately 75	Approximately 10	Approximately 300	500 (I-D-S/I user base)
CONFIGURATION				
CPU's supported	Any Burroughs B 6700 or B 7700	Univac 90 /30 , 90/60, 90/70, Series 70 (Virtual)	Any Univac 1100 Series	Honeywell Series 60 600,6000
Operating systems	MCP	OS/3; VS/9	1100 Operating System (EXEC 8)	GCOS
Minimum memory requirements	180K bytes	90/30—131K bytes; larger systems—262K bytes	15K words plus buffer space	12K words plus buffers
DATA BASE FEATURES				
Data base organization	Link, network, ring, hierarchical	Network, hierarchical	Network, hierarchical	Hierarchical, network
Application languages	COBOL, ALGOL, PL/1	COBOL, FORTRAN , Assembler	COBOL, FORTRAN, Assembler	COBOL 74
Data base languages	DASDL	DDL, DML	DDL, DML	Data Definition Language
Access methods supported	All MCP-supported access methods	CODASYL access methods	CODASYL access methods	Random, sequential, index
Variable-length segments	Yes	No	Yes	No
Data base security	MCP password capability	Logical record locks, passwords	System level file security	Password, privacy keys to field level
System accounting facilities	System Log	Yes	Yes	Summary of utilization statistics
RECOVERY FEATURES				
Checkpoint/restart	Yes	Yes	Yes	Yes
Data base integrity	Audit trail and transaction logging	Before/after imaging	Audit trail, before/after imaging	Roll-back and automatic recovery
OTHER SYSTEM FEATURES				
Concurrent batch/on-line	Yes	Planned for release 4.1	Yes	Yes
Concurrent application program	Yes	Planned for release 4.1	Yes	Yes
Inquiry/retrieval facility	INQUIRY Language	UNIQUE, planned for release 3.1	Query Language Processor	Interactive I-D-S/II, Data Query System
Report Generator	Reporter System	None	COBOL Report Writer	MDQS (planned for 1977)
Data dictionary support	None	User-defined	Yes	Yes
Telecommunication interfaces	NDL and MCS	None	Yes	TPS
PRICE (basic system, no options)	$24,000; monthly plans available	No charge to Univac customers	No charge to Univac customers	$692/month (includes COBOL 74)
COMMENTS	DMS-II is integrated to some degree with the MCP control program. DMS-II facilities are incorporated into Burroughs higher-level languages	DMS/90 is a CODASYL-oriented DBMS. It interfaces with the operating system for many DBMS facilities. Many enhancements are being implemented	DMS 1100 is a CODASYL-oriented DBMS. It uses many of the operating system data management facilities, and offers highly flexible networking features	I-D-S/II is the latest version of the Honeywell DBMS. Conversion capabilities for I-D-S/I users will be available in 3rd quarter 1976. I-D-S/II is a CODASYL-type DBMS

A Buyer's Guide to Data Base Management Systems

COMPARING THE LEADING DATA BASE MANAGEMENT SYSTEMS (Continued)

System	IDMS	IMS	INQUIRE
DATAPRO 70 report number	70E-272-02	70E-491-01	70E-498-01
Vendor	Cullinane Corp.	IBM	Infodata Systems, Inc.
Current number of users	120	Datapro estimates about 500	Approximately 70
CONFIGURATION			
CPU's supported	S/360, 370; Univac Series 70	System/360, 370	S/360 Model 40 and up, S/370 Model 135 and up
Operating systems	DOS, DOS/VS, OS, OS/VS, EDOS; TDOS, VMOS	IMS-2—OS, OS/VS; IMS/VS—OS/VS	OS, OS/VS, CMS
Minimum memory requirements	55K bytes plus 10K for each batch job; 2K for on-line jobs	IMS-2—128K bytes and up (DB); 512-768K (DB/DC); IMS/VS—90K and 350K	130K—Command Language; 40K—Procedural Language interface
DATA BASE FEATURES			
Data base organization	Hierarchical, network	Hierarchical (sequential and direct)	Hierarchical, network, relational
Application languages	COBOL, PL/1, ASM, FORTRAN, RPG II	COBOL, PL/1, Assembler	COBOL, FORTRAN, PL/1, Assembler
Data base languages	DDL, DNL	Data Language/1	Inquire Command Language
Access methods supported	BDAM	VSAM, ISAM, OSAM, BSAM, VTAM, SAM	ISAM, VSAM, BDAM
Variable-length segments	Yes	Yes, with VSAM	Yes
Data base security	Password protection and subschema	Password and terminal access limitation	Encryption and password protection
System accounting facilities	Automatic logging of system statistics	System Log Analysis tape and utilities	Data base for usage accounting routines
RECOVERY FEATURES			
Checkpoint/restart	Utilities supplied; automatic with TP	With IMS/DC only	None
Data base integrity	Via prohibitive access	Transaction backout logging	Backout and image logging
OTHER SYSTEM FEATURES			
Concurrent batch/on-line	Yes	Yes	Yes
Concurrent application program access	Yes	Yes, in DB/DC mode	Yes
Inquiry/retrieval facility	No; one is planned for delivery in 1976	IQF, GIS/VS	INQUIRE Command Language
Report generator	CULPRIT, EDP/AUDITOR	GIS/VS, GIS-2	Command Language
Data dictionary support	User-defined	An FDP is available for IMS-2 users	User-defined
Telecommunication interfaces	Most standard TP monitors	CICS and IMS/DC	CICS, TSO, IMS/DC, CMS, and others
PRICE (basic system, no options)	$40,000; rental and lease also available	Monthly license: IMS-2—$616/mo.; IMS/VS—$770/mo.	$39,500; monthly lease available
COMMENTS	CODASYL-type DBMS. Compression feature and user entry points. Forms Approach retrieval due in 1976	Up to 255 segment types per logical record with 15 levels. Fine-tuning utilities included	Data reference by multiple keys. Multi-Data-base Processor feature

A Buyer's Guide to Data Base Management Systems

COMPARING THE LEADING DATA BASE MANAGEMENT SYSTEMS (Continued)

System	MODEL 204	SYSTEM 2000	TOTAL
DATAPRO 70 report number	70E-174-01	70E-652-01	70E-132-01
Vendor	Computer Corp. of America	MRI Systems Corp.	Cincom Systems, Inc.
Current number of users	22	Over 100	Approximately 1000
CONFIGURATION			
CPU's supported	System/360, 370	S/360, 370; Univac 1100; 1100; CDC 6000 and Cyber Series	Most major mainframes; also Varian, DEC PDP-11, and IBM S/3 minicomputers
Operating systems	OS, OS/VS	OS, OS/VS, CMS; EXEC 8; SCOPE, KRONOS, NOS	All associated operating systems
Minimum memory requirements	120K bytes plus buffer space	140-200K bytes– IBM; 32K words—Univac; 20K words—CDC	40K bytes on S/360; 8-10K bytes on minis
DATA BASE FEATURES			
Data base organization	Hierarchical, network	Hierarchical, network	Network
Application languages	COBOL, FORTRAN, PL/1, Assembler	FORTRAN, COBOL, PL/1, Assembler	COBOL, FORTRAN, PL/1, Assembler, RPG II
Data base languages	IFAM/II	DDL, IMMEDIATE	DBDL, DML
Access methods supported	QSAM, EXCP level	All standard IBM access methods	BDAM, DAM
Variable-length segments	Yes	Yes	Physical—no; logical—yes
Data base security	Password lockout, log-in protection	Password lockout, assigned authority	None
System accounting facilities	Multi-user accounting log and utilities	Logs, statistics, and estimation tools	None
RECOVERY FEATURES			
Checkpoint/restart	Yes	Yes	Yes (with TP monitor)
Data base integrity	Rollback and audit trail	Transaction log and activity audit	Logging, dump and restore
OTHER SYSTEM FEATURES			
Concurrent batch/on-line	Yes	Yes	Yes
Concurrent application program access	Yes	Yes	Yes
Inquiry/retrieval facility	User Language	System 2000 Query/ Update facility	For Honeywell systems only
Report generator	User Language	Yes	SOCRATES
Data dictionary support	User-defined	DDL	None
Telecommunication interfaces	CICS, Intercomm, and self-contained DC	TP 2000, CICS, TSO, Intercomm	ENVIRON/1, CICS, TASK/ MASTER, Intercomm
PRICE (basic system, no options	$83,200; rental and lease also available	$30,000 plus $10,000 for language interface	IBM DOS—$34,500; IBM OS—$39,500
COMMENTS	Supports 250 physical files which can be cross-referenced by a single user. Multi-threadng and data independence	Can handle up to 9 strings simultaneously with Multi-Thread option	A CODASYL-Type DBMS. Supports up to 32 levels of data elements and up to 65,000 files

TOTAL
Cincom Systems, Inc.

MANAGEMENT SUMMARY

With approximately 1000 installations at this writing, TOTAL is by far the most widely used independently vended data base management system. It is also available for operation on a much wider range of computer equipment than any competitive DBMS, thereby offering users with diverse hardware installations an opportunity to standardize their data base software.

TOTAL is a host-language data base management system implemented much along the lines of the CODASYL Data Base Task Group Report, except that the user can use other host languages as well as COBOL. TOTAL provides an effective means for organizing and managing diverse data to make it both efficient and convenient for application programmers to maintain and retrieve the data for processing. TOTAL performs in both batch and on-line environments, with Cincom's own ENVIRON/1 (Report 70E-132-02) or PMI's Intercomm (Report 70E-694-01) available to handle the data communications and SOCRATES available for data retrieval and display.

As compared with IBM's IMS (Report 70E-491-01), TOTAL offers a more simplified and streamlined approach to data base management, providing a more satisfactory solution for the majority of users with data base applications and utilizing far fewer resources in terms of personnel and hardware. Of course, it is incumbent upon each user to ensure that he is, in fact, within that majority. IMS also has a growing following, and has many strengths of its own. But very few of us need a Mack truck to pick up our groceries.

TOTAL can manage virtually an unlimited number of data sets on an "integrated, non-redundant" basis and provides for association of each of these data sets with other data sets to form an integrated data base. TOTAL allows the user to relate data across many functional and/or departmental boundaries, permitting the data processing applications to mirror the system of management within an organization.

Basically, TOTAL is oriented toward simplifying application programming problems by handling complex relationships among the various data items that constitute the information base in a user's shop. Such complex relationships are quite common for most installations where multiple but related applications are implemented. One of the very nice conceptual features of TOTAL is that users can start small, with perhaps a single application using two or three data files, and grow in modular fashion to very large and sophisticated data base/data communications structures with very little impact on the previously implemented systems.

TOTAL's ability to handle complex relationships among data items stems from the fact that TOTAL enables the user to utilize any logical data structure (network or hierarchical) in conjunction with a single physical data organization. In network structures, a data record may have any number of subordinate or "member" records.

TOTAL, the leader among independently marketed data base management systems, provides facilities for data base generation and accessing by any host language that supports a CALL statement. It is operational on IBM System/360 and 370, IBM System/3, Control Data, Honeywell, NCR, Univac, and Varian computer systems.

CHARACTERISTICS

SUPPLIER: Cincom Systems, Inc., 2300 Montana Avenue, Cincinnati, Ohio 45211. Telephone (513) 662-2300.

BASIC FUNCTION: TOTAL provides facilities for generation of a data base that permits automatic cross-referencing among data records. A facility is also provided for accessing the data base from conventional application programs written in COBOL, PL/1, FORTRAN, etc.

OPERATION: The TOTAL Data Base Management System is composed of several phases, one of which is also called TOTAL. The basic system includes three phases: one is for generating the program for controlling the data base structure, one is for pre-formatting the disk areas, and one is for controlling the access to the data base.

Phases of TOTAL are dynamically loaded at execution time, permitting data read/write/add/delete or structural maintenance, reading data only, and generating master record addresses only (which speeds the loading of the data base).

TOTAL permits the establishment of two types of records: a single-entry or master record and a variable-entry record. Each group of records, of either type, forms a file (data set). Linkages can be set up that permit automatic retrieval of all variable-entry records associated with a particular single-entry record based on the linkage. A variable-entry record can be part of many linkage paths or chains.

A TOTAL data base is composed of multiple data sets or files. Linkages can exist between any master file and any variable-entry file. Multiple file data bases can be established. A particular master file or variable-entry file can be part of more than one data base. The multiple paths of access allowed by such a structure, called a network structure, simplify the logic of application programs using the data. In the case of TOTAL, they also reduce the amount of disk storage required to hold information by eliminating duplicate fields or records.

To one familiar with sequential and hierarchical sequential files, the benefits of a network structure are not immediately evident. It seems at first glance that the power of a network structure is limited because only one sublevel of linking is possible; i.e., master to variable-entry. The real power of this structure lies in the fact that multiple master files can be established, each for a particular relationship, and any number of variable files can be related to any number of these master files. Each variable file can handle any number of different record types.

The three topics relevant to discussing what TOTAL can do for you are data base generation (DBGN), data base definition language (DBDL), and data management language (DML).

TOTAL
Cincom Systems, Inc.

But, more importantly, each member record may also have any number of "owner" records. In hierarchical structures (such as IBM's IMS), an owner record may have many members—but each member can have only a limited number of owners.

The network structure technique of TOTAL allows the access paths to data to be different from those on the physical storage device(s). The system uses two types of data sets called "single entry" and "variable entry." Access to single-entry or master records is direct, according to the value of a control field, (e.g., customer number, employee number, part number, etc.). Access to a variable-entry record is normally through a particular master record and a particular relationship, which is called a linkage path. There may be any number of paths to access variable-entry records. Because relationships among data items can be specified directly, related items can be retrieved directly.

TOTAL provides a bidirectional hierarchical structure technique which allows for logical development of any number of parent, child, sub-child, etc., relationships. The user simply defines the desired relationship utilizing the single-entry master data set in conjunction with the variable data set.

TOTAL provides a simple and straightforward set of facilities for organizing a data base and for manipulating the data in the data base within an application program. To the application programmer, who can be coding in COBOL, FORTRAN, PL/1, Assembly Language, or any other host language that supports a CALL statement, accessing information in the data base is much like using a subroutine to calculate and return data values; file definitions and data structure definitions are not required for the information obtained from or written into the data base. All data base access from the host language is through CALL statements.

TOTAL provides a complete logging, backup, and recovery capability. When used with ENVIRON/1, automatic warm or cold restarts in real-time and data communications environments are available. This facility, combined with periodic data base dumps, allows full recovery from any physical file loss.

TOTAL's dynamic data base logging (which captures "before" images of data fields modified by transactions) allows automatic restoration of the data base to its original form or to the most recent checkpoint in system processing. "After" image, function, and transaction logging are also provided.

One serious problem of data base implementation in general is the impact of change. If an item of data is used across a broad range of applications, then both the data processing impact and the corporate implications of change to the data or its structure can be significant in two areas. First, every program that uses the data item being changed will require, at a minimum, recompilation. Second, the data base and data structures will often be exposed to serious reorganization problems. This can involve the shifting of internal corporate responsibilities among user groups as well as large expenditures of data processing resources.

DATA BASE GENERATION: The DBGN program accepts the data base structural definitions to DBDL and outputs an assembly-language program. This program contains the definition of the fields within each record, the data sets and their relationships to one another, and the buffers. Core requirements for this module are typically 500 bytes per data set plus buffers. After assembly, this module is catalogued to be core-resident with the application program. After executing the FORMAT program, which formats disk storage, application programs can then be run using the TOTAL module (DML) to access the data base. The first application program is usually one that loads data into the data base; this same program can then be used to add additional records to a file after initial creation. Output from DBGN is a listing of the data base definition statements which serve as a guide to the contents and structure of the data base records.

A randomizing algorithm is used to calculate master record physical addresses based on the value of the control field. If duplicate addresses are calculated, a pointer is used in that record to show where the "duplicate" or synonym record is stored. Thus, the complete disk space allocated can be used. Once all space is used up, the data file must be reloaded with new parameters. Cincom provides a utility to handle such occurrences.

DATA BASE DEFINITION LANGUAGE: Writing the control statements for structuring the data base is relatively simple. A short preamble identifies the data base by name and is followed by a definition of the data elements (fields) contained in each single-entry and variable-entry data set. The relationships (linkages) between data sets are also defined, along with the physical attributes of each data set (device type, total number of records, and blocks per track).

A part of the description of each data set is the name of the I/O buffer, the names of all linkage fields, and the structure of the records. TOTAL processes only fixed-length records; i.e., all records in a data set must be of the same length. Only one format can be specified in single-entry data sets, but multiple formats are allowed in variable-entry data sets, each identified by a two-digit record code. Linkages between a single-entry data set and a variable-entry data set can be qualified by the record code, if desired.

Data element structure is identified in a very straightforward manner by stating the name and size in bytes. No data format identification is required or can be specified. The data names refer to data elements, the smallest chunk of data that can be named for retrieval. In many versions of TOTAL, data elements can be subdefined. The data elements may or may not refer to individual fields used in an application program.

Comments can be easily incorporated in the data definition statements and are highly recommended to serve as aids to application programmers.

DATA MANAGEMENT LANGUAGE: DML provides the facilities for retrieving information and passing it to an application program, as well as for opening and closing files and for adding and deleting records from an existing data set. Absent from DML are any provisions for establishing new data sets or establishing new linkages among existing data sets. Some users consider this to be an advantage since it does not allow application programmers to bypass the data base design and generation phase and possibly compromise the existing data base structure.

DML functions at the CALL level in the host programming language. For data transfer operations, the parameter list includes an operation code, the name of the file (data set) to be used, the name of a field to hold the status returned at the end of each operation, the name of the master file record control field, the name of the I/O area, and a list of the data elements to be transferred. If variable-entry files are being accessed, the linkage path and field are also identified. For some operations, such as opening or closing a file, not all parameters are required. The operation specified controls the parameters needed.

TOTAL
Cincom Systems, Inc.

MANAGEMENT SUMMARY

With approximately 1000 installations at this writing, TOTAL is by far the most widely used independently vended data base management system. It is also available for operation on a much wider range of computer equipment than any competitive DBMS, thereby offering users with diverse hardware installations an opportunity to standardize their data base software.

TOTAL is a host-language data base management system implemented much along the lines of the CODASYL Data Base Task Group Report, except that the user can use other host languages as well as COBOL. TOTAL provides an effective means for organizing and managing diverse data to make it both efficient and convenient for application programmers to maintain and retrieve the data for processing. TOTAL performs in both batch and on-line environments, with Cincom's own ENVIRON/1 (Report 70E-132-02) or PMI's Intercomm (Report 70E-694-01) available to handle the data communications and SOCRATES available for data retrieval and display.

As compared with IBM's IMS (Report 70E-491-01), TOTAL offers a more simplified and streamlined approach to data base management, providing a more satisfactory solution for the majority of users with data base applications and utilizing far fewer resources in terms of personnel and hardware. Of course, it is incumbent upon each user to ensure that he is, in fact, within that majority. IMS also has a growing following, and has many strengths of its own. But very few of us need a Mack truck to pick up our groceries.

TOTAL can manage virtually an unlimited number of data sets on an "integrated, non-redundant" basis and provides for association of each of these data sets with other data sets to form an integrated data base. TOTAL allows the user to relate data across many functional and/or departmental boundaries, permitting the data processing applications to mirror the system of management within an organization.

Basically, TOTAL is oriented toward simplifying application programming problems by handling complex relationships among the various data items that constitute the information base in a user's shop. Such complex relationships are quite common for most installations where multiple but related applications are implemented. One of the very nice conceptual features of TOTAL is that users can start small, with perhaps a single application using two or three data files, and grow in modular fashion to very large and sophisticated data base/data communications structures with very little impact on the previously implemented systems.

TOTAL's ability to handle complex relationships among data items stems from the fact that TOTAL enables the user to utilize any logical data structure (network or hierarchical) in conjunction with a single physical data organization. In network structures, a data record may have any number of subordinate or "member" records.

TOTAL, the leader among independently marketed data base management systems, provides facilities for data base generation and accessing by any host language that supports a CALL statement. It is operational on IBM System/360 and 370, IBM System/3, Control Data, Honeywell, NCR, Univac, and Varian computer systems.

CHARACTERISTICS

SUPPLIER: Cincom Systems, Inc., 2300 Montana Avenue, Cincinnati, Ohio 45211. Telephone (513) 662-2300.

BASIC FUNCTION: TOTAL provides facilities for generation of a data base that permits automatic cross-referencing among data records. A facility is also provided for accessing the data base from conventional application programs written in COBOL, PL/1, FORTRAN, etc.

OPERATION: The TOTAL Data Base Management System is composed of several phases, one of which is also called TOTAL. The basic system includes three phases: one is for generating the program for controlling the data base structure, one is for pre-formatting the disk areas, and one is for controlling the access to the data base.

Phases of TOTAL are dynamically loaded at execution time, permitting data read/write/add/delete or structural maintenance, reading data only, and generating master record addresses only (which speeds the loading of the data base).

TOTAL permits the establishment of two types of records: a single-entry or master record and a variable-entry record. Each group of records, of either type, forms a file (data set). Linkages can be set up that permit automatic retrieval of all variable-entry records associated with a particular single-entry record based on the linkage. A variable-entry record can be part of many linkage paths or chains.

A TOTAL data base is composed of multiple data sets or files. Linkages can exist between any master file and any variable-entry file. Multiple file data bases can be established. A particular master file or variable-entry file can be part of more than one data base. The multiple paths of access allowed by such a structure, called a network structure, simplify the logic of application programs using the data. In the case of TOTAL, they also reduce the amount of disk storage required to hold information by eliminating duplicate fields or records.

To one familiar with sequential and hierarchical sequential files, the benefits of a network structure are not immediately evident. It seems at first glance that the power of a network structure is limited because only one sublevel of linking is possible; i.e., master to variable-entry. The real power of this structure lies in the fact that multiple master files can be established, each for a particular relationship, and any number of variable files can be related to any number of these master files. Each variable file can handle any number of different record types.

The three topics relevant to discussing what TOTAL can do for you are data base generation (DBGN), data base definition language (DBDL), and data management language (DML).

TOTAL
Cincom Systems, Inc.

While the impact on corporate organization is essentially a management problem (and one common approach to solving this problem is through establishment of a Data Base Administrator function), TOTAL handles both of the EDP-related considerations very nicely by providing data independence at the data field or data element level. New fields of data can be added to a physical record with no impact upon operational programs which use fields contained in the old record. TOTAL features integrated data base facilities, but with separate data files. This means that data within one physical data file can be logically related to data in a different physical data file. This approach minimizes the impact of physical change to the data base.

There are two principal differences between TOTAL and monolithic data base management systems such as IMS. First is TOTAL's use of one physical data structure (direct) to provide all required logical data structuring (network or hierarchical) and all required processing needs (transaction handling, batch updating, and reporting). This capability simplifies the efforts of the data base designer and provides for improved performance in many application areas.

The second difference is also, perhaps, TOTAL's principal limitation. The defined data base structure is fixed, not dynamic, once the data base has been loaded. Records can be added to and deleted from any of the existing data sets. But if new relationships (linkage paths) are to be defined, new data sets cannot be added and disk storage areas cannot be expanded without at least partial regeneration of the data base, which may entail reloading of all affected data files.

However, TOTAL is one of the most effective high-powered software systems in use today, and it has been successful in displacing IBM's IMS in numerous installations. In September 1972, Honeywell Information Systems announced an agreement to market the Honeywell Series 200/2000 version of TOTAL through its own worldwide sales force. In addition, Cincom has developed TOTAL for use on the NCR Century Series, the CDC Cyber Series, the UNIVAC Series 70 and 9400/9700, the IBM System/3, and the Varian V 70 computers, assuring further recognition and market acceptance for the already widely used TOTAL system.

A comparison of TOTAL with other current data base management systems can be found in Report 70E-010-61.

USER REACTION

In the 1975 Datapro survey of proprietary software users, a whopping 116 users responded with their ratings of TOTAL. That number of respondents was exceeded by only two other software products (Panvalet and CICS). What's more, TOTAL's users rated it among the very best of the data base management systems currently on the market. TOTAL came within 0.2 point of making the Software Honor Roll, with a 3.3 rating in overall satisfaction. In general, the major data base management systems were not rated as highly as most of the other types of software considered in the survey.

The exact procedures vary depending on the host programming language. The chief difference among the various languages supported is whether or not literals can be used in the parameter list.

A total of five sets of functions is provided: one for opening and closing files individually or via a list, one for serial retrieval of records from a single- or variable-entry file, one for resetting the serial record counters, one for working with single-entry (master) files, and one for working with variable-entry files. There are also control functions which allow the dynamic loading of TOTAL with the proper data base definition at execution time, and control functions which specify the type of logging desired.

Serial processing refers to taking the records in the physical sequence in which they appear on disk storage, not sequentially according to the value of a control field. Variable-entry records can be processed according to their actual layout on disk or according to any linkage path. Normally this would mean that there is some order to the records. Serial processing provides some efficiencies when all records in a data set must be handled.

Processing functions available for master files include reading, writing, addition, and deletion of a master record. All variable records linked to a master record must be deleted in a separate operation before a master record can be deleted. The space occupied by a deleted record is available immediately for reuse by TOTAL.

A host of operations is available for working with a variable-entry data set. All references to a variable-entry data set are along a linkage path with a specific master record. Records can be read forward or in reverse (i.e., starting from either the beginning or end of a chain), or a relative location can be specified for retrieval. Much flexibility is provided for adding a variable-entry record to a chain. Records can be added at the beginning or end of a chain or inserted within the chain. In this way, a logical sequence can be maintained. A delete function is also included.

The use of the CALL subroutine method makes accessing data very much like the calculation of values in a subroutine. Instead of performing a calculation, TOTAL retrieves data. No file and data specifications are necessary in the host program other than definition of the logical record or fields that the host program will manipulate. Because only a data element list is required in the CALL statement, modification of the physical data record structure to include new fields does not require that application programs referencing the affected files be changed.

MULTIPROGRAMMING: The multi-tasking versions of TOTAL can operate in either or both of two modes. The LOCAL mode allows the servicing of a single application task or multiple application sub-tasks within a single partition or region. The CENTRAL mode services any number of application tasks or sub-tasks residing in other partitions or regions while TOTAL resides in an independent partition. Single-tasking, or batch, TOTAL can be used in a multi-tasking environment. In this case, however, individual copies of TOTAL must appear in each partition or region. Any number of TOTALs may coexist in the machine, regardless of version or mode.

Under TOTAL, a given application can process either the entire data base or any portion desired by defining a subset of files and requesting that definition by name at execution time. The sub-definition identifies that application task's "area of influence." Having specified a definition at execution time, the application may not access any portion of the data base not described in this definition. Any number of data base definitions can be constructed and selected at execution time, but a given application task can select only one at a time.

Under the single-task version of TOTAL, only one application task may update (add, delete, write) a given

TOTAL
Cincom Systems, Inc.

In overall satisfaction, ease of installation, and ease of use, TOTAL was rated excellent by well over one-third of the respondents, and in every one of Datapro's rating categories at least 64 percent of the respondents rated the product either excellent or good. The only negative aspects of the package revealed by our survey were the facts that 28 of the 116 respondents indicated that they considered it costly and 40 rated the documentation no better than fair. On the other hand, 88 respondents indicated that the package was flexible and 86 noted that TOTAL saves programming time.

The respondents had been using TOTAL for an average of 21 months, so our survey represents a group of users who had a significant amount of time to evaluate the evolution of the system. TOTAL's main memory utilization for these 116 respondents averaged 40K bytes–an impressively low figure for a full-scale DBMS. Virtually no real criticism was reported against the package, and this highly positive endorsement by more than 10 percent of TOTAL's current users is a real vote of confidence in the package as well as in its vendor.

The summarized user ratings were as follows:

	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	51	53	12	0	3.3
Throughput/efficiency	29	68	13	2	3.1
Ease of installation	47	48	10	2	3.3
Ease of use	44	58	9	0	3.3
Documentation	11	61	34	6	2.7
Vendor technical support	31	53	23	5	3.0
Training	18	62	22	3	2.9

*Weighted Average on a scale of 4.0 for Excellent.

Based on the above ratings and the number of satisfied customers of TOTAL, it is fair to say that no one who is evaluating data base management systems should omit TOTAL from the list of prime candidates.□

file at a time. This is controlled by the file reservation facility invoked by the task when the file is logically opened. However, under multi-tasking versions, application tasks may share files by selecting the same data base definitions at execution time. Concurrent accessing and updating are allowed for commonly used files.

To protect a data base from concurrent updating of a record, TOTAL provides a record "holding" facility. When a record is read for a specific task, TOTAL will hold the logical record for that task until the task issues a write or issues another request to the same file. Any application task attempting to access a record currently held by another task is prohibited from processing that record until it has been released. Subsequent attempts to access the record will cause the application task to "stall." TOTAL will monitor the occurrence of these attempts and will cause a release of the record, allowing the application task to proceed. A status indication will be returned to the original owner of the record, indicating that the record has been released and the transaction must be reprocessed. Application tasks which will not need to update records can indicate that fact to the system; any attempt to update a record will then be refused by TOTAL.

Data base definition under the multi-task version of TOTAL allows any number of buffers to be specified for any I/O area desired and any number of files within that definition to share any I/O area.

HARDWARE/SOFTWARE REQUIREMENTS: TOTAL is currently operational on IBM System/360 and 370 computers running under DOS, DOS/VS, OS, VS1, and VS2; Honeywell Series 200 and 2000 computers running under Mod 1 (MSR), Mod 2, and OS/2000; UNIVAC Series 70 and 9400/9700 computers running under TDOS and DOS; NCR Century Series; IBM System/3; and CDC Cyber Series computers. A version is also available for use on Varian V 70 Series minicomputers. Host languages include COBOL, FORTRAN, PL/1, RPG II, and Assembly Language. TOTAL "CENTRAL" is available for most of the above systems. Main memory requirements range from a minimum of 3K bytes for the read-only TOTAL to 30K bytes for the full "CENTRAL" version. These figures exclude application code and I/O buffers.

PRICING: Prices vary, depending on the target computer system. TOTAL is available on either a purchase or rental basis. For IBM DOS and DOS/VS installations, the base purchase price is $34,500 and the rental price is $825/month. For IBM OS, OS/VS, and corresponding CICS installations, the base figures are $39,500 purchase and $1,150/month rental. The Intercomm interface, which is usable only under OS and OS/VS, can be purchased for $4,250 or leased for $200/month. Several versions of TOTAL are available for IBM computers that may be priced slightly higher, depending on the options and utilities chosen. System engineering and education support are bundled for the first 60 days of product usage.

INITIAL DELIVERY: Early 1969.

CURRENT USERS: About 1000 as of January 1976.■

IDMS
Cullinane Corporation

MANAGEMENT SUMMARY

The popularity of data base management systems is increasing rapidly, but there is much confusion in the minds of potential users as to the types of systems they should be implementing. As with programming languages, there are many users and vendors alike who advocate *standardization* of data base management concepts and techniques. As a result of standardization, COBOL has become a programming language that is practically universal in the business data processing world, leading to a degree of machine and manufacturer independence that was virtually undreamed of just a decade ago. A CODASYL committee task force has developed a comparable set of language specifications for data bases, in the hope that their specifications will become the basic standards for all data base management systems.

Currently, users spend a great deal of time and money in planning for data base management systems. Once a particular system has been acquired and the user has fully committed himself to its implementation, he is not likely to convert to another system for a long time to come; the expense of conversion would probably be prohibitive. So it is essential that the original decision regarding the right DBMS for that installation be completely compatible with the company's immediate needs and future growth plans, and that the system complement the user's hardware configuration, both at present and in the future. It is important that the DBMS vendor be responsible and in a position to react to future hardware innovations. It is also important that the data base management system be flexible enough to adapt easily to application and system software changes as they occur.

Cullinane's IDMS (Integrated Database Management System) satisfies these requirements and occupies an important position in the data base software market: it is one of the few existing data base systems that has been specifically designed to meet current CODASYL language specifications and is available for IBM System/360 and 370 computers.

IDMS was originally developed in a "Fortune 500" company in 1970 and 1971, and was put into production there in early 1972. In 1973, Cullinane was awarded complete responsibility for the system, including all technical developments, enhancements, field support, and marketing. Since then, Cullinane has produced four major releases of IDMS, and a fifth is scheduled for December 1977. Enhancements have generally followed the CODASYL guidelines, with emphasis on performance and usability.

IDMS represents a highly flexible approach to application system development that permits the designer to effectively model his application requirements without technical complexity. The system reportedly allows for installation in 30 minutes, and, according to the vendor, customer personnel are trained on-site to use IDMS in three days. While IDMS boasts features absent from most competitive systems, Cullinane states that it requires a very modest allocation of system resources in compari- ➢

IDMS, a 1976 Datapro Software Honor Roll package, is a DBMS based on the current CODASYL data base system specifications. The system is operational on IBM System/360 and 370 computers running under DOS, OS, and their VS counterparts. It comes complete with a built-in data dictionary, a full set of utilities, an On-Line Query language, and interfaces to numerous teleprocessing monitors.

CHARACTERISTICS

SUPPLIER: Cullinane Corporation, 20 William Street, Wellesley, Massachusetts 02181. Telephone (617) 237-6601. Technical support locations include New York, Cincinnati, Pittsburgh, Chicago, Washington, Atlanta, and San Francisco.

BASIC FUNCTION: IDMS is a data base management system that is operational on IBM System/360 and 370 computers running under DOS, OS, or their VS counterparts. It conforms with the CODASYL Committee Specifications, including the 1973 Data Definition Language report and the 1975 Data Manipulation Language report.

The complete system encompasses a data base design methodology, a language to describe the physical and logical data base (DDL), data manipulation language (DML) compilers for COBOL and PL/1, and a data base manager that provides record storage, control, space management, security, and backup and recovery functions. Also included is a data base dictionary subsystem. IDMS operates in local mode (in which each user program invokes its own copy of the data base management code), as well as in central mode (in which one copy of the data base management system services any number of user programs—batch and/or on-line—concurrently). Interfaces allow it to be used with most of the commercially available telecommunications monitor systems, including the Cullinane-marketed Shadow II system and IBM's CICS.

OPERATION: By using the graphic language and design methodology, the system designer or data base administrator (DBA) constructs a pictorial view of the data base system. Then, using the schema data definition language (DDL), the system designer defines the data base globally. Next, using the device media control language (DMCL), the designer establishes necessary mappings between the data base and actual secondary files. Finally, using the subschema DDL, the designer describes one or more data base "views" which can be used by application programs. The above-named descriptions are stored by IDMS in a central Data Directory file. This directory provides an informational source from which IDMS produces Data Dictionary reports for use by the data administrative staff, end users, and applications analysts and programmers.

The data base as defined globally in the schema consists of "data items," "records," "sets," and "areas." Items are individual fields of data, records are aggregates of related items handled as a unit, sets are relationships between records, and areas are regions of the IDMS data space. The schema compiler accepts, validates, and stores this description of the data base into the Directory file.

A subschema—or logical—view of the data base as it appears to application programs consists of one or more designated areas, one or more set relationships, and portions ➢

IDMS
Cullinane Corporation

son with competitive systems, yet offers greater performance. Because of the method used in the architecture of IDMS, new applications can be implemented with minimum involvement of production systems, and current applications can grow in volume without affecting performance ratios. IDMS data bases frequently do not require periodic maintenance or reorganization. Data base recovery software is a standard feature of the system.

A comparison of IDMS with other current data base management systems can be found in Report 70E-010-61.

USER REACTION

IDMS was one of three DBMS products that earned places on the 1976 Datapro Software Honor Roll, marking the first time in the four-year history of the Datapro's annual survey of proprietary software users that any DBMS made the Honor Roll. Seventeen respondents rated IDMS, and the average rating for each of the seven categories was 3.0 or higher. In fact, the overall satisfaction category was almost unanimously rated as "excellent." Here's how these 17 users rated the system:

	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	14	3	0	0	3.8
Throughput/efficiency	6	9	1	0	3.3
Ease of installation	9	5	2	0	3.4
Ease of use	10	5	1	0	3.6
Documentation	4	9	4	0	3.0
Vendor technical support	10	5	1	0	3.6
Training	9	8	0	0	3.5

*Weighted Average on a scale of 4.0 for Excellent.

There weren't many comments on the survey forms submitted by the users, but it is fair to say that the ratings themselves tell a rather conclusive story regarding the users' feelings towards IDMS. If one had to pick some area where the product could be improved, it would have to be documentation, which four users rated only "fair." Even so, the overall rating of 3.0 for the documentation was well above the industry average.

Of the 17 respondents, 14 indicated that IDMS performed as advertised "immediately" and the other 3 said it did so "eventually." Eleven of the users reported that no modifications were required to install the system, while five said that the vendor performed some modifications. One user did not respond to this question. Only one respondent indicated that any additional cost was involved in the modifications, and that figure was negligible compared to the overall cost of the system.

The average length of user experience with IDMS was 12.3 months, and the average main storage used was 82.5K bytes. Of the 17 respondents, 10 had IDMS installed on IBM System/370 computers, 4 were using System/360's, and 2 had it installed on Univac Series 70 computers. One user did not identify the system upon which IDMS was installed. The operating systems supporting IDMS included DOS/VS, OS, and the various versions of OS/VS. One user had Software Pursuits' DOS/MVT replacement operating system installed on his System/360 and was using IDMS very

of records which contain the data needed for a particular application. Meaningful names for the above items are expressed in 1 to 16 characters. Privacy locks are also applied to areas, sets, records, and functions at the subschema level. The subschema compiler accepts, validates, and stores this decription of the application data base into the Directory as well.

The IDMS concept of area usage provides a technique to isolate a portion of the data base for a particular application, thereby limiting the extent of the data base required to be on-line for that production run.

By using the set concepts, IDMS satisfies the data base management concept of the interrelationship of record types. Within a set, one record type functions as the "owner" and one or more record types function as "members." Using the set concept, hierarchical, network, partially inverted, indexed, and bill of material data bases can easily be defined. Set characteristics are defined by the system designer and consist of independent choices of set order, set membership, and set linkage.

The designer can select one of five logical orders for each set:

- SORTED—members are stored under control of a logical sort field.
- FIRST—members are stored LIFO (last-in, first-out).
- LAST—members are stored FIFO (first-in, first-out).
- NEXT—members are stored in a descending sequence under control of the application program.
- PRIOR—members are stored in ascending sequence under control of the application program.

The same member record may be in a different sequence in each set in which it participates.

The designer can select one of four membership specifications for each set. The choices are:

- Mandatory Automatic—members are connected automatically into a set at the time they are stored and remain in the set until erased from the data base.
- Mandatory Manual—members are connected into a set under program control but remain in the set until erased from the data base.
- Optional Automatic—members are connected automatically into a set when stored but may be disconnected from one set and connected to another under program control.
- Optional Manual—members are connected into sets as well as disconnected from sets under program control.

A member record may have different linkage specifications in each set in which it participates. Four linkage options are available for each set:

- NEXT—the system maintains unidirectional pointers for processing in the forward direction only.
- NEXT and PRIOR—the system maintains bidirectional pointers for processing in forward as well as reverse order.
- NEXT and OWNER—the system maintains pointers back to the respective owner in each member record as well as pointers in the forward direction.
- NEXT, PRIOR, and OWNER—a combination of the second and third options noted above.

IDMS
Cullinane Corporation

successfully. Several of these users had IDMS interfaced with well-known teleprocessing monitor systems.

When asked to indicate the major advantages of the system, 15 respondents noted that IDMS saves human resources, 14 said it was flexible, 7 mentioned that it saved system resources, and 2 cited the system as inexpensive. On the negative side, two users said that the system was complex, one said it was costly, and one reported that some compatibility problems were encountered initially and that the system was slow. In the case of IDMS, the users clearly feel that the advantages far outweigh the drawbacks.

As mentioned previously, 1976 was the first year in which any data base management systems were elected to the Datapro Software Honor Roll. The fact that IDMS was one of the three members of this elite group certainly warrants its evaluation if you're in the market for a flexible DBMS that conforms to CODASYL specifications.□

Records are stored into the data base by one of the following three techniques:

- CALC—provides for record storage based on a primary key within the data record used by the data base system to calculate a relative storage address. CALC is used to define entry points into the data base and is often used for "master" type data. Duplicate keys may be accepted or rejected based on design criteria.
- VIA—provides for storage of member records physically near the owner record to which they are related within the set. VIA is often used for "transaction" type data and provides for more efficient processing because all the associated data are brought into main memory with a single access.
- DIRECT—the application program directs the data base management system to store a record at a given relative location. This technique is useful when the application program desires to establish a custom addressing scheme.

In all cases, actual record storage, space management, buffering, and control are the responsibility of the data base management system.

The subschema provides for a logical subset of the data base and defines the rules by which the individual application system may access the data base. Data independence and data security features are implemented in the subschema, and the user may define additional privacy control through the use of passwords, special keys, or data range analysis. IDMS also provides special routines for data compression and decompression, variable-length records, editing and validation, record substitution, auditing, statistical analysis, and encoding/decoding through the use of special DBA procedures.

Application programs access the data base through the use of the DML, which provides the interface to the data base system. DML commands such as STORE, ERASE, CONNECT, DISCONNECT, MODIFY, and OBTAIN obviate the need for coding CALL statements. The DML commands and the host-language statements, which may be used in COBOL or PL/1 (each has its own syntax so the DML commands are a natural extension of the host language) are read into a DML processor, which checks the syntax of each command, ensures its logical consistency with the data base, and inspects the security and privacy locks associated with the application subschema.

In addition, the DML compiler builds the user data work areas, data base communication blocks, and data base declarative statements.

Assembler-language programmers access the data base through the use of a set of IDMS DML macros. FORTRAN users can communicate with the data base system through the use of CALL statements.

IDMS includes a comprehensive collection of data base utilities which provide for:

- Formatting and loading.
- Dumping and restoration.
- Recovery from log file (forward and backward).
- Restructuring.
- Display of selected sections of the data base.
- Statistics on program and space usage.

Options to the basic system consist of:

- Central Version, which permits one copy of the data base manager to service multiple batch regions and on-line application programs.
- CULPRIT, the vendor's output processor and retrieval system (Report 70E-272-03).
- Integrated Data Dictionary, a full-capability incorporated dictionary intended for the control of all types of data—data base as well as conventionally organized. The IDD is built around the Data Directory file discussed above.
- Telecommunications monitor interfaces, including one for Shadow II, a multi-tasking, multi-threading system marketed by the vendor.

HARDWARE/SOFTWARE REQUIREMENTS: IDMS operates on any IBM System/360 or 370 computer under DOS, EDOS, DOS/VS, OS, VS1, or VS2. The basic system requires approximately 55K bytes of main storage. Each batch application requires approximately 10K bytes, and each on-line application requires 2K bytes of memory. All storage requirements may be satisfied with virtual storage. An adaptation of IDMS called DBMS-11 is being marketed by Digital Equipment Corporation for use on the DEC PDP-11 minicomputer.

PRICING: IDMS is available under license, lease, and rental plans. The one-time license fee is $42,000 for IDMS, $15,000 for the Central Version, $20,000 for CULPRIT, $9,000 for the On-Line Query, and $12,000 for the Integrated Data Dictionary. The license applies to all computers at any one site. Additional sites may be included at 65 percent of the initial license price. An annual usage fee of 10 percent applies to all sites and entitles the user to all new system releases, new documentation, technical support, program maintenance, membership in the IDMS User Association, and newsletters and technical bulletins.

A corporate-wide license is available which permits unrestricted use of IDMS at all corporate computing locations. IDMS is available to government users under GSA scheduling, and educational institutions may qualify for an educational discount.

INITIAL DELIVERY: Cullinane began marketing IDMS in May 1973.

CURRENT USERS: 175 as of January 1977.■

IMS
IBM Corporation

MANAGEMENT SUMMARY

IBM's Information Management System (IMS) is one of the most controversial and frequently discussed software systems on the market today. It represents the criterion by which other vendors judge or compare their data base management systems, and it has introduced numerous innovative data base management concepts that have been copied, and frequently improved upon, by other vendors. It is always easier to improve upon a new concept than it is to come up with the concept in the first place. This is not to say that other DBMS vendors have not implemented new concepts; but, by and large, IMS has been leading the way.

As a consequence of its leadership role in the complex realm of data base management, IMS is probably the most misunderstood and misused software product on the market, and has been called everything from the world's most sophisticated and technologically current software system to a core-eating, time-consuming monstrosity.

One of the primary assets of IMS is its flexibility. To be sure, IMS is one of the most flexible systems on the market today, having more options and "bells and whistles" available to the prospective user than almost any other system available. But this type of flexibility carries along with it the complexity of creation and implementation, not to mention use.

IMS, then, is a flexible and complex system with many subtle attributes that are capable of causing considerable impact on any organization that implements it. It is therefore imperative that prior to implementing IMS, or any DBMS for that matter, the preliminary work-up of system requirements on the part of the user be as complete as possible. IMS is not the best answer for all types of user environments, and users that invest in IMS only to find that it doesn't meet their needs will be hard-pressed to convert to another system without incurring a great deal of expense and lost time. On the other hand, users who can make good use of IMS will find that this thorough preliminary planning will take much of the mystery out of the system.

Briefly, IMS is concerned with the efficient organization and structure of data items on physical direct-access devices. It provides a means for physical-level access to the data, and sets up an interface between the user's application program and the operating system's data management and communications management facilities.

The fundamental functional concept underlying IMS's technique for data base management is fairly simple. The user employs utility programs to describe the structure of the system from two viewpoints: stored data structure as seen by the system, and logical data structure as seen by an application. Stored data is described only once, but many descriptions of the logical data can exist. The logical data base descriptions are external to the application programs and exist as stored data which the system references when processing

IMS, IBM's principal data base management system, is available in several versions that permit operation in either real or virtual storage mode and in either batch or on-line environments. This report describes and analyzes IMS-2, IMS/VS, and three IMS subsets for use on smaller systems: DL/1 DOS/VS, DL/1 Entry, and VANDL-1.

CHARACTERISTICS

SUPPLIER: IBM Corporation, 1133 Westchester Avenue, White Plains, New York 10604. Telephone (914) 696-1900.

BASIC FUNCTION: IMS provides the capabilities for generating and accessing a data base, with automatic cross-referencing among data records. IMS/VS operates under the OS/VS1 or OS/VS2 operating system. IMS-2 (IMS/360) provides the same basic functions as IMS/VS (although IMS/VS contains many more capabilities) and operates under OS/MFT or OS/MVT—or, with fixed pages, under OS/VS1 or OS/VS2. Both IMS/VS and IMS-2 offer on-line message processing with the optional Data Communications (DC) Feature. With this feature, on-line inquiry with IQF (Interactive Query Facility) or GIS/VS (General Information System) and batch inquiry with GIS or GIS/VS are available. In addition, a data language (DL/1), whose function is to replace user I/O coding with simpler commands to IMS, is provided.

The same basic data base facilities are provided for DOS/VS users with DL/1 DOS/VS. However, its only available data communications extension is linkage to CICS/DOS/VS (Report 70E-491-02). IMS/VS can also link to CICS/VS, and IMS-2 to CICS/OS Standard, but they must then forego their own data communications features and, more significantly, IQF.

IBM has also recently released DL/1 Entry, which became available in July 1975. It is a compatible subset of DL/1 DOS/VS and is upward-compatible from another IBM data base management system, VANDL-1. VANDL-1 is now a class C product, having been downgraded from class B when DL/1 Entry was announced. DL/1 Entry also supports a communications interface to CICS/VS.

All of the foregoing programs are written in Assembly language and offer their DB facilities to users of COBOL, PL/1, and Assembly language.

OPERATION: IMS operation requires that the user perform the following steps (presented in the typical development sequence, although considerable overlapping normally takes place):

- The applications must be defined in terms of functional requirements and types of data to be processed. Individual applications programs must then be written in COBOL, PL/1, or Assembly language to perform these applications.

- All of the individual data requirements for each application must be coordinated into an overall data base contents requirement. This task is most appropriately handled by an individual serving as a Data Base Administrator. Based upon the processing logic of the various applications programs, the administrator then selects one of the hierarchical IMS structures for the physical storage of the data.

access requests to the data base from application programs. Thus, IMS maps logical data descriptions into the physical data descriptions it alone can use, using logical data names supplied by the application. It also determines an access strategy and performs the requested function against the stored data. The entire range of systems attributes and functions relies upon this basic concept.

A major benefit of the centralized data base approach is what IBM promotes as "data integrity." In the case of IMS, this means having the management and planning staffs all "reading off the same scorecard"; i.e., one source of figures is used by everyone. Surprisingly (to the uninitiated), this basic coordination is very difficult to achieve in most large corporations.

THE VERSIONS OF IMS

IMS was originally developed as a joint venture between IBM and North American Rockwell Company in the mid-1960's. At that time, the classic concept of special-purpose file processing programs with one or more individual files for each program was the accepted way of doing business. In an effort to centralize file maintenance, eliminate the storing and maintenance of redundant data, isolate individual programs from each other and from the data (data independence), ease the often complex handling of variable-length records, and provide the framework for applications program development, NAR and IBM produced IMS-1. The system was subsequently released as a Program Product with full IBM support in September 1968.

Although IMS passed through a number of intermediate releases before arriving at Version 2, the major distinctions of IMS-2 from IMS-1 can be briefly summed up as: the addition of two hierarchical direct data organizations–HDAM and HIDAM; the inclusion of Interactive Query Facility (IQF) into IMS-2 only (this did not occur simultaneously with the initial release of IMS-2); the significant improvement of performance features within IMS-2; and the different packaging of IMS-2, with data communications handled as an optional feature.

Early in 1973 the next major step in IMS evolution occurred when IBM announced IMS/VS (IMS for virtual storage), available under OS/VS1 and OS/VS2. The initial release was made available in February 1974, and a version for OS/VS2 with multiprocessing followed in July 1974.

This spelled the trend for all future enhancements: the fact that all enhancements would be to the VS version of IBM's data base management product. In April 1975, Version 1.1 was released by IBM. The major improvements can be summarized as follows:

- Data base changes can be "backed out" in the event of an abnormal termination of a message or batch-message program.
- Batch checkpoint/restart facilities were added for batch and batch-message processing.

- Communications and other on-line retrieval requirements must be evaluated, both in terms of functional (application) needs and in terms of the network design criteria (terminal types and quantities) that may be present. If communication-type requirements or interactive query requirements are present, a choice must be made of two data communications monitors (the Teleprocessing Option of IMS or CICS) and/or two high-level inquiry methods (IQF or GIS). Note, however, that CICS and IQF cannot both be selected: IQF runs only with the IMS Data Communications feature.
- At this point, the Data Base Administrator must put together the IMS modules required to support his installation. The complexity of this task can be considerable, particularly if data communications is involved. The IMS modules include the resident nucleus, control facility system tasks, BTAM, GAM, OSAM, or ISAM re-entrant modules, BTAM or BSAM device support, terminal handlers, control blocks, buffer pools, and certain OS requirements.
- In conjunction with the selection of an appropriate data storage structure and data communications or inquiry method, the Data Base Administrator then sets out to create one or more "logical" data bases that redefine the actual physical data base. This process establishes logical relationships in the form of pointers and/or chains that facilitate user-program access to the data. Although wide flexibility is provided within IMS to create valid logical data bases that tie in various data elements or segments, inappropriate design at this point can result in severe overhead problems.
- The actual data base is created by a special-purpose off-line IMS utility program.
- Finally, the user is ready to incorporate DL/1 access language into his own programs and process his IMS data base applications. This step basically consists of replacing the standard I/O syntax in each application program with CALL statements containing parameters pertinent to IMS.

The above sequence of operational steps describes a classical, theoretical implementation of the full-scale IMS management information system environment, including data communications capability. In common practice, however, most IMS applications evolve incrementally, with only a basic core of applications and a partially developed data base to start with. Also, any number of responsibilities assigned to the Data Base Administrator above can be handled by systems programmers at the user's option.

The data base environment within IMS provides for a "separation" between the data and each user program. This concept of data independence implies that changes to the data base–such as the inclusion of new fields, changes in record length or description, and physical reorganization into new structures or different device types–need not be accompanied by corresponding changes in what is likely to be a large number of individual application programs that access the data base.

In order to achieve data independence, a common symbolic program linkage and data base description is used that supports five primary types of data-base access or I/O operations:

- Retrieve a unique segment (GET UNIQUE).
- Retrieve the next sequential segment (GET NEXT).
- Replace the data in an existing segment (REPLACE).
- Delete the data in an existing segment (DELETE).
- Insert a new segment (INSERT).

These operations can be performed on one or more hierarchically related segments or data elements in a

IMS
IBM Corporation

RELATIONSHIPS BETWEEN THE IMS VERSIONS AND OPERATING SYSTEMS AND COMMUNICATIONS PROGRAMS

Program ID	Operating System (and mode)				Interface for:
	DOS	DOS/VS	OS	OS/VS	
IMS Version 2; 5734-XX6	NA	NA	yes	real	DC feature or CICS Version 1 or 2
IMS/VS; 5740-XX2	NA	NA	NA	virtual	DC feature or CICS/OS/VS
DL/1 DOS/VS; 5746-XX1	NA	virtual	NA	NA	CICS/DOS/VS
VANDL-1; 5799-AEY	yes	virtual	NA	NA	CICS DOS or CICS DOSE
DL/1 Entry; 5746-XX7	NA	virtual	NA	NA	CICS/DOS/VS

- New program isolation allows the same segment types within data bases to be concurrently updated by multiple-message and batch-message programs.
- Improved multiple indexing allows access to root and dependent segments within a data base, allowing any field within most segments within a data base on disk storage to be indexed.
- VSAM support is provided for disk storage of data; VSAM data sets and single-element HISAM data bases share a common stored record format.
- The IMS control region and dependent regions can be executed in virtual mode.
- Variable-length data segments can be dynamically changed.
- Support for remote computers includes the System/3 and System/7, both on nonswitched multipoint lines.

IMS/VS communications now supports all devices supported under the communications feature of IMS-2 except for the 1030 terminal. Also, the IMS/VS control region must operate in real mode when the 7770 Model 3 Audio Response Unit is used.

Additional enhancements were made to IMS/VS in Release 1.1.1 in July 1975. The most important of these include:

- VTAM support.
- Support for the 3600 Financial System and the 3790 Communication Systems.
- Improvements to the Message Format Services (MFS) function.
- Enhancements to systems definition (SYSGEN).
- Support for the 3740 Data Entry System.
- The incorporation of a Utility Control Facility (UCF) which provides for utility operations and maintenance in preparation for recovery and restart under the control of one job, one step, and in one scheduling.

logical data base record. Each of the operations is invoked by a Data Language/1 (DL/1) command. (The appropriate DL/1 command is indicated in parentheses above.) DL/1 works with a data base description (DBD) produced by the off-line IMS utility program that creates the data base. The DBD provides the "mapping" from the logical structure of the data base (as viewed by the application program) to the physical structure of the data base (as kept on a storage device by OS). The logical data structure of IMS is based upon segments: an IMS data base consists of 1 to n data base records; a data base record consists of 1 to n segments of up to 255 segment types and up to 15 segment levels. There is one root segment per data base record and 0 to n occurrences of dependent segments per parent.

Four primary physical data organizations are provided in IMS:

- Hierarchical Sequential Access Method (HSAM)–an extension of basic serial tape and disk file processing (SAM). This method offers limited data independence and no interrelatability of the data base through "pointers." In order to insert a data base record, the data base must be copied up to that point, the new record written, and the rest of the data base copied. Each record is physically present in the serial order in which it logically appears in the data base.
- Hierarchical Indexed Sequential Access Method (HISAM)–provides an imbedded hierarchy of ISAM-like data sets that are related by sets of symbolic pointers or keys. The distinguishing aspect of HISAM (or HSAM), as opposed to the hierarchical direct methods described below, is that all segments in a physical data base record are "related by physical juxtaposition." For HISAM, this means that a direct-access relationship is established between all of the physical blocks containing the segments belonging to a given data base record. An Overflow Sequential Access Method (OSAM) physically contains the segments that cannot fit in the HISAM logical record. OSAM is based upon standard OS physical data sets and combines the best features of both BSAM and BDAM: concurrent sequential and direct access for retrieval, in-place updating as well as addition at the end of a data set, data set "end" recognition, secondary extent definition for data sets, etc. IMS/VS also provides support for certain VSAM data sets on disk; that is, VSAM data sets and single-segment HISAM data bases will share a common stored record format. HISAM does not yield particularly good results in an on-line environment.
- Hierarchical Direct Access Method (HDAM)–stores data in a physical tree structure with all segments in a physical data base record related by direct addresses. Segments can be interrelated to each other as physical twins (multiple occurrences of the same segment type under a given parent), physical parents (segment

- An IMS/VS DB Monitor which provides a trace of internal activities of the system and a report program that prints summary and distribution reports of this collected data.

- The General Sequential Access Method (GSAM) which allows batch programs to access BSAM and VSAM ESDS data sets through DL/1, and provides data set repositioning through Extended Checkpoint/Restart (XRST).

- Additional support for the System/7 attached on a binary synchronous, nonswitched, contention or polled line, including program load/IPL capability.

DOS and DOS/VS users have not been left out of the IMS picture. In October 1972, IBM announced DL/1 DOS/VS, which was subsequently delivered in November 1973. This program—not to be confused with DL/1 (Data Language/1), the command language used with IMS—provides DOS/VS users with data base capabilities that are a compatible subset of those found in the DB portion of IMS/VS. For DOS/VS users interested in a DB/DC system, DL/1 DOS/VS offers an interface to CICS/DOS/VS (Report 70E-491-02).

IBM released DL/1 Entry in March 1975. This new product, intended for small-to-medium scale installations running under DOS/VS, is a compatible subset of DL/1 DOS/VS and is upward-compatible from VANDL-1. The only access methods DL/1 Entry supports are HISAM and HSAM, but it does support applications programs written in RPG II as well as COBOL, PL/1, and Assembler language.

VANDL-1, a programming RPQ item, was the only available DB facility designed for use by DOS or DOS/VS users prior to the release of DL/1 DOS/VS and DL/1 Entry. As a result of the introduction of these two products and the conversion aids developed to allow VANDL-1 users to upgrade, there are few VANDL-1 users still around. In fact, IBM downgraded the service classification of VANDL-1 to class C in April 1975.

THE STRUCTURE OF IMS

IMS-2 and IMS/VS each come in two basic versions: Data Base (DB) system only, and Data Base system with Data Communications feature (DB/DC). While the DB system handles input job streams by batch scheduling, the DB/DC system is transaction-oriented and schedules work based upon input messages.

In the DB batch environment, all necessary IMS modules are combined with each user's program in their individual partitions. The DB/DC system, on the other hand, concurrently supports both batch and on-line applications and provides an independent control partition or region for the IMS modules that is separate from the message processing/batch applications processing partitions or regions, each of which contains only the user's application program.

IMS is a "host language" system, which means that no direct application code is provided by IMS. The user writes his own applications programs in COBOL, PL/1, immediately above), or physical children (first and last occurrence of each segment type immediately subordinate) through chains of pointers. HDAM uses OSAM as a base for data storage and provides very effective access to dependent segments—especially in teleprocessing environments—at some overhead cost in terms of data base size.

- Hierarchical Indexed Direct Access Method (HIDAM)—provides an ISAM index to data physically stored in OSAM format. The ISAM index contains the key of a root segment and a direct address to the root segment, while the actual storage of data is done in OSAM data sets. Because the data base index and the actual data base are kept on two separate data sets, reorganization of the index separately from the data is facilitated. HIDAM is the most generally appropriate and most often used data organization method for IMS applications.

In addition to the above data structures and access methods, the basic batch-oriented version of IMS (also called "DL/1 Data Base" or the DB system), can be augmented with data communications capability to produce a transaction-driven system. The DB system is prerequisite to the DC Feature ("IMS teleprocessing"). The resulting full-scale IMS is known as the DB/DC system, and can handle both batch and on-line operations concurrently. A DB/DC system can have a wide variety of physical terminals, each of which can have one or more logical or symbolic names. Individual security parameters can be associated with each terminal's logical name. Among other facilities in a DB/DC system are the following:

- Master Terminal for network control, message scheduling, and interrogating or altering the IMS processing functions.
- Input message processing for transactions, terminal-to-terminal messages, or system interface messages.
- System log with a queue of all I/O messages on disc (used for restart).
- Standardized message editing with the ability to add user-defined editing routines.
- Conversational processing capability for terminal access to applications programs.
- Video Terminal Paging feature.
- System security both by terminal and by password.
- Terminal Test Mode to debug on-line IMS applications.
- Message scheduling, including up to 15 priority levels, message class, and dynamic message re-prioritization parameters.
- Support for the 3740 Data Entry System.
- Utility Control Facility.
- DB and DC Monitor for resource utilization activity recording and printing.
- 2740, 2741, and 3600 support under MFS.
- Support for the 3790 for limited data base inquiry and response. (This feature is not supported under IQF.)
- VTAM network sharing support.
- Multiple Physical Page Input Messages support from multiple 3270 pages.

As an alternative to the IMS Teleprocessing option, a DB/DC system can be put together using the Customer Information Control System. CICS generally provides similar functional capabilities with lower overhead in

IMS
IBM Corporation

- An IMS/VS DB Monitor which provides a trace of internal activities of the system and a report program that prints summary and distribution reports of this collected data.

- The General Sequential Access Method (GSAM) which allows batch programs to access BSAM and VSAM ESDS data sets through DL/1, and provides data set repositioning through Extended Checkpoint/Restart (XRST).

- Additional support for the System/7 attached on a binary synchronous, nonswitched, contention or polled line, including program load/IPL capability.

DOS and DOS/VS users have not been left out of the IMS picture. In October 1972, IBM announced DL/1 DOS/VS, which was subsequently delivered in November 1973. This program—not to be confused with DL/1 (Data Language/1), the command language used with IMS—provides DOS/VS users with data base capabilities that are a compatible subset of those found in the DB portion of IMS/VS. For DOS/VS users interested in a DB/DC system, DL/1 DOS/VS offers an interface to CICS/DOS/VS (Report 70E-491-02).

IBM released DL/1 Entry in March 1975. This new product, intended for small-to-medium scale installations running under DOS/VS, is a compatible subset of DL/1 DOS/VS and is upward-compatible from VANDL-1. The only access methods DL/1 Entry supports are HISAM and HSAM, but it does support applications programs written in RPG II as well as COBOL, PL/1, and Assembler language.

VANDL-1, a programming RPQ item, was the only available DB facility designed for use by DOS or DOS/VS users prior to the release of DL/1 DOS/VS and DL/1 Entry. As a result of the introduction of these two products and the conversion aids developed to allow VANDL-1 users to upgrade, there are few VANDL-1 users still around. In fact, IBM downgraded the service classification of VANDL-1 to class C in April 1975.

THE STRUCTURE OF IMS

IMS-2 and IMS/VS each come in two basic versions: Data Base (DB) system only, and Data Base system with Data Communications feature (DB/DC). While the DB system handles input job streams by batch scheduling, the DB/DC system is transaction-oriented and schedules work based upon input messages.

In the DB batch environment, all necessary IMS modules are combined with each user's program in their individual partitions. The DB/DC system, on the other hand, concurrently supports both batch and on-line applications and provides an independent control partition or region for the IMS modules that is separate from the message processing/batch applications processing partitions or regions, each of which contains only the user's application program.

IMS is a "host language" system, which means that no direct application code is provided by IMS. The user writes his own applications programs in COBOL, PL/1,

immediately above), or physical children (first and last occurrence of each segment type immediately subordinate) through chains of pointers. HDAM uses OSAM as a base for data storage and provides very effective access to dependent segments—especially in teleprocessing environments—at some overhead cost in terms of data base size.

- Hierarchical Indexed Direct Access Method (HIDAM)—provides an ISAM index to data physically stored in OSAM format. The ISAM index contains the key of a root segment and a direct address to the root segment, while the actual storage of data is done in OSAM data sets. Because the data base index and the actual data base are kept on two separate data sets, reorganization of the index separately from the data is facilitated. HIDAM is the most generally appropriate and most often used data organization method for IMS applications.

In addition to the above data structures and access methods, the basic batch-oriented version of IMS (also called "DL/1 Data Base" or the DB system), can be augmented with data communications capability to produce a transaction-driven system. The DB system is prerequisite to the DC Feature ("IMS teleprocessing"). The resulting full-scale IMS is known as the DB/DC system, and can handle both batch and on-line operations concurrently. A DB/DC system can have a wide variety of physical terminals, each of which can have one or more logical or symbolic names. Individual security parameters can be associated with each terminal's logical name. Among other facilities in a DB/DC system are the following:

- Master Terminal for network control, message scheduling, and interrogating or altering the IMS processing functions.
- Input message processing for transactions, terminal-to-terminal messages, or system interface messages.
- System log with a queue of all I/O messages on disc (used for restart).
- Standardized message editing with the ability to add user-defined editing routines.
- Conversational processing capability for terminal access to applications programs.
- Video Terminal Paging feature.
- System security both by terminal and by password.
- Terminal Test Mode to debug on-line IMS applications.
- Message scheduling, including up to 15 priority levels, message class, and dynamic message re-prioritization parameters.
- Support for the 3740 Data Entry System.
- Utility Control Facility.
- DB and DC Monitor for resource utilization activity recording and printing.
- 2740, 2741, and 3600 support under MFS.
- Support for the 3790 for limited data base inquiry and response. (This feature is not supported under IQF.)
- VTAM network sharing support.
- Multiple Physical Page Input Messages support from multiple 3270 pages.

As an alternative to the IMS Teleprocessing option, a DB/DC system can be put together using the Customer Information Control System. CICS generally provides similar functional capabilities with lower overhead in

IMS
IBM Corporation

or Assembly Language, for which interface modules are provided. These application programs access the data base through Data Language/1 (DL/1) commands that replace what would otherwise be complex procedural I/O coding and the file descriptions present in applications programs. Thus, most applications programs written for an IMS environment tend to be smaller than similar non-IMS programs that perform the same function. For COBOL, as an example, the source language programs tend to be one-fourth to one-third smaller, and the corresponding object programs up to one-fourth shorter, with a consequent reduction in application program development effort.

Each IMS "physical" data base is organized in one of four hierarchical or tree-type structures. The user views the physical data through a "logical" data base description that allows sophisticated security precautions and data relationships to be expressed. All user interfacing is done at the "segment" level, which is the basic IMS data element. The logical data base represents one of IMS's best assets, as it facilitates data independence and simplifies applications program logic. Data stored in multiple separate physical data bases can be treated as it it were logically part of one data base. Because of the logical data base concept and DL/1, existing IMS applications programs can be insensitive to the physical reorganization of data on the storage devices, the addition of new applications or data, changes or new developments to OS or VS access methods, and changes in device types or terminal devices.

A full complement of IMS utility programs is provided to describe the data base structure, create the data base, reorganize the data base, recover and reconstruct data (checkpoint/restart capability), specify security control, and analyze the system workload.

One interesting aspect of IMS is its potential for management information system (MIS) development. MIS and IMS get confused with each other for obvious reasons; but an MIS system can be thought of as a collection of management information application programs depending upon an IMS data base. The extent to which these application programs satisfy management's information needs determines the extent to which the MIS can be considered successful. Often the installation of IMS is the preliminary step in the implementation of a full-scale MIS.

CORPORATE IMPLICATIONS

Perhaps more important than any of the computer-related considerations of a data base management system are the overall corporate implications. As an example, consider the case where different departments access the same data in the data base. Which department is to be responsible for gathering the update information, editing it, and entering it as maintenance transactions to the IMS data base? Further, if the department currently responsible for the legwork and procedures involved in updating some of the common elements in the data base receives a new charter eliminating that department's need to access that data, can it drop the update responsibility? Who would then inherit the update manpower and/or responsibility? Since there may not be a change in the net overall corporate responsibilities, it would

some environments. CICS was designed for relatively short program modules of about 2K to 6K bytes, while the IMS Teleprocessing option is better suited to 20K-byte modules or larger. The CICS/IMS interface has been available since early 1972.

Also available for full DB/DC IMS systems is the Interactive Query Facility (IQF). IQF is a basic query language that offers the capability for on-line retrieval and display of data in an IMS data base. IQF consists of retrieval phrases that define, delete, list, sort, count, total, limit, and query the data base; the "when" qualifier to establish criteria for data selection; and a basic complement of relational (EQ, NE, LT, GT, LE, GE), logical (AND, OR), and arithmetic (+, -, /, *) operators. Other handy features of IQF include null words (e.g., THE, OF, FOR) set up by "define" phrases, literal and numeric constants, segment synonyms, etc.

IQF is handled like a standard applications program under IMS. A batch utility program is used to create and maintain the following data bases associated with the use of IQF:

- System Data Base for rapid resolution of data base field names so that each IQF query does not have to specify individual segment names.
- Phrase Data Base that contains predefined phrases and null words set up by the user to tailor IQF for a specific set of applications.
- Two IQF Index (QINDEX) Data Bases that index the non-key data in the IMS data base(s). Typically, a "large" and a "small" key-field data base are created. The QINDEX data bases are also used to streamline on-line retrieval requests.

The IQF utility is run when putting the IMS system together and for subsequent index creation or modification. IQF can be used only with the IMS DC Feature. It is not supported by CICS. IQF can be used in conjunction with the more powerful query capability of the Generalized Information System (GIS/2)–a full-scale information system–although no direct relationship between IQF and GIS/2 exists.

GIS/2 and GIS/VS, with the DL/1 Query Support Feature, can be used to produce tailored processing modules that permit batch and on-line operations upon IMS data bases as well as a variety of other data file types. GIS is actually a sort of "super RPG" that accepts report format and query selection criteria as input and produces an object deck plus Job Control Language (JCL) as output. These GIS object modules can be executed under control as batch programs only. Full GIS query capabilities (except LIST RECORD and HOLD RECORD) are supported for IMS data bases.

A Bill Processor Bridge System is also available that converts Bill of Material Processor (BMP 360A-ME-06X) and Data Base Organization and Maintenance Processor (DBOMP 5736-XX4) files into IMS data bases.

With Releases 1.1 and 1.1.1, some additional features have been incorporated into IMS. These enhancements and additions can be summarized as follows:

- Support for the IBM System/7, attached on a binary synchronous, nonswitched, contention or polled line.
- The Utility Control Facility, which provides the means of performing most data base utility operations and maintenance in preparation for recovery and reorganization under control of one job, one step, and in one scheduling. It handles data base recovery and data base manipulation in reorganization, combining of data base changes into composite change records in change accumulation, and backup copy processing, all in semi-automatic fashion.
- A Data Base Monitor, which provides the ability to trace internal activities of an IMS/VS DB system and includes a report program that prints summary and

IMS
IBM Corporation

seem that no net manpower change should be effected; but individual departments can be dramatically impacted. And, of course, there is a matter of security. Who gets to see the payroll information, since it may possibly be on the same physical storage device as the corporate production data?

There are direct solutions to some of these problems, but not to others. Top corporate involvement in the design and subsequent care and feeding of the data base is a key element to success with IMS. One approach that has gained almost complete acceptance is to establish a Data Base Administrator function. This position calls for a sort of "referee" to control the contents of the data base, take charge of enforcing adherence to system standards, and steer potential intraorganizational conflicts over the data base to the proper authorities for resolution. This staff function must be neutral to all parties and technically qualified to give advice to management. The Administrator has many responsibilities akin to those of the system programmers, and one of the most important of these is "tuning" the system for maximum performance. This can involve reorganization, creation, or modification of logical data base descriptions, etc.

An important trade-off that must be made by the designers of the data base and subsequently reviewed by the Data Base Administrator during the tuning process is that of DL/1 programming complexity and data organization method (HIDAM, HSAM, etc.) against processing time, size of disk storage requirements for the data base, and main memory requirements for the IMS system modules. Adjustments will be called for as the data base develops, and tuning should be an on-going process with periodic reviews.

USER REACTION

In the 1975 Datapro survey of proprietary software users, we received 35 responses from IMS users and 32 responses from DL/1 users. These users rated the systems as follows:

IMS	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	5	20	8	2	2.8
Throughput/efficiency	2	20	5	6	2.5
Ease of installation	3	14	9	8	2.4
Ease of use	2	17	9	5	2.5
Documentation	5	22	6	0	3.0
Vendor technical support	8	17	8	1	2.9
Training	4	15	7	1	2.8

DL/1	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	4	18	7	3	2.7
Throughput/efficiency	3	12	12	2	2.6
Ease of installation	6	7	12	4	2.5
Ease of use	5	11	10	3	2.6
Documentation	4	11	11	4	2.5
Vendor technical support	6	9	9	5	2.6
Training	3	10	8	4	2.5

*Weighted Average on a scale of 4.0 for Excellent.

Both systems rated higher in the 1975 survey than they did in our 1974 survey, although the differences were not dramatic except perhaps in the area of documentation. In 1974 IMS was rated 2.6 in documentation, whereas in 1975 the rating went up to 3.0. DL/1

distribution reports of this collected data. The Monitor offers the user the ability to accumulate performance data on specific application designs, data base designs, and resource allocation.

- The General Sequential Access Method (GSAM), which allows batch programs to access OS BSAM and VSAM ESDS data sets through DL/1. GSAM provides data set repositioning through Extended Checkpoint/Restart (XRST), another release enhancement.

- Parallel DL/1, which permits calls to data bases to be processed in each message processing region with the exception of calls that require reading from an ISAM data set, which must be done in the IMS/VS control region. This eliminates the necessity of routing all DL/1 calls through the control region. In a multiprocessing environment, it permits DL/1 execution to occur in both CPU's simultaneously.

HARDWARE/SOFTWARE REQUIREMENTS: IMS/VS can be used on System/370 Models 145, 155-II, 158, 165-II, and 168. In addition to the user's application programs, IMS/VS itself requires a minimum region or partition size of 90K for DB only or 350K (Release 1.0 and 1.1) for DB/DC. When it is incorporated with the user's application programs, practical minimum main storage requirements are 384K bytes for OS/VS1, 512K for OS/VS2 Release 1, and 768K for OS/VS2 Release 2 or 3. With the Data Communications feature added, these systems would require approximately 512K, 768K, and 1024K bytes, respectively. Also required are an OS/VS1 or OS/VS2 system console, at least one 2400 or 3400 series 9-track tape unit, and direct-access space for system libraries and working storage. With 2316 packs, the DB system requires 125 cylinders for Version 1.0 and 138 cylinders for Version 1.1, and the DB/DC system requires 225 cylinders for Version 1.0 and 248 cylinders for Version 1.1. With 3336 packs, a minimum of 83 cylinders (Version 1.0) or 91 cylinders (Version 1.1) is required for the DB system and 158 cylinders (Version 1.) or 174 cylinders (Version 1.1) for the DB/DC system. Also, one of the following nonswitched devices is required to support the IMS/VS Master Terminal: 2740 Communications Terminal with Station Control feature, 1050 Data Communications System with 1052 Printer-Keyboard, or 3270 Information Display System.

IMS-2 operates on System/360 and 370 computers with a minimum of 128K bytes for DB and 256K for DB/DC under MFT. Under MVT, the system requires 256K for batch-only DB and 512K for DB/DC. Under OS/VS, minimum real storage requirements for DB are 384K for OS/VS1 and 512K for OS/VS2. For DB/DC, the requirements are 512K and 768K, respectively. In addition, the DB system requires one 7- or 9-track magnetic tape unit and 740 cylinders of 2311 disk storage (or equivalent) for program storage and work space. The DB/DC system requires an additional non-switched 1050 Data Communication System or 2740 Communication Terminal and two 9-track tape units.

In addition to the above requirements, IQF requires a message processing region of 50K. A batch IMS region of at least 200K is required to run the IQF utility. This is based on a SORT work area size of 44K. If a larger work area size was specified at SYSGEN time, an increase must be made to the minimum region size for the IQF utility.

DL/1 DOS/VS requires a System/370 Model 125, 135, or 145 with a minimum of 98K bytes for a batch system or 163K with CICS. The practical minimum number of 2K pages for batch DL/1 is 30 with all functions except logical relationships and secondary indexing, and 40 for all functions. With CICS, these figures are 45 pages and 70 to 100 pages, respectively. Also required are a DOS/VS system console, two 9-track tape units, and 75 cylinders of 2316 disk storage space (or equivalent) for a batch system. For a CICS system, 150 cylinders of 2316 (or equivalent) disk storage space are required.

DL/1 Entry, a subset of DL/1 DOS/VS, runs on a System/370 Model 115 or larger under DOS/VS and

IMS
IBM Corporation

documentation ratings also went from 2.2 in 1974 to 2.5 in 1975. Overall satisfaction remained the same for IMS but climbed from 2.4 to 2.7 for DL/1. The other categories either gained a tenth of a point or lost the same amount in the latest Datapro survey.

Responses received from the users indicated that the major advantages were flexibility and the saving of programming time. The IMS counts were: flexible–22, and saves programming time–17. The DL/1 breakdown was: flexible–16, and saves programming time–17.

The principal disadvantages cited were as follows: for IMS, 13 respondents criticized the product as being too costly, 21 noted it as too complex, and 16 said it used excessive resources. DL/1 users complained that the product was too complex (11 respondents) and that it used excessive resources (15 respondents).

In response to the question "Did the product perform as advertised?", 11 IMS users answered Immediately, 18 responded Eventually, and 1 answered Never. For DL/1 the responses were: 13 Immediately, 8 Eventually, and 3 Never.

Another significant question addressed to the users was "Did the system require modifications?" To this the responses were as follows: for IMS, 19 answered No; 6 answered Yes, by the vendor; and 6 more answered Yes, by the user. For DL/1, 18 answered No; 6 answered Yes, by the vendor; and 5 answered Yes, by the user. While the average modification cost cited by the IMS users was only $200, the DL/1 users reported an average cost of $8,929 for modifications. This difference probably attests to the fact that since IMS has so many built-in facilities, the user has only to select those required to meet his needs, while the DL/1 user is more limited in terms of what the system will support.

In the Datapro user survey, the average number of months in use for IMS was 14, while that for the DL/1 product was 10 months. The average main storage usage reported for IMS was 265K bytes, and for DL/1 it was 93K bytes. The majority of users responding to the survey were Virtual Storage system users, which attests to IBM's wisdom in concentrating on enhancing the VS versions of the data base management system.

Although the 1975 survey of proprietary software users was conducted strictly through a mailing of questionnaires, we have had the added advantage of speaking by telephone with numerous users of IMS and DL/1, and there were many sound comments made during our conversations.

Several users, in commenting about the complexity and cost of installing IMS, noted that they had profited by adding one or more people to their staffs who had gained prior experience with the IBM DBMS. This move saved the users many hours of grief and also enabled them to define more accurately the resources that would be needed to bring the system up. This is critically important advice in our opinion. We have noted before that IMS is a highly flexible and complex system, but with proper training and the right kind of personnel, the transition to IMS can be made with a minimum of anxiety.

requires at least a 96K processor. It affords the user a CICS interface along with many of the features of DL/1 DOS/VS. DL/1 Entry is a superset of, and upward-compatible with, VANDL-1.

VANDL-1, now a class C RPQ software product, runs on either a System/360 or 370 under DOS or DOS/VS. VANDL-1 requires at least 14K bytes of main memory for the system plus an additional 8K to 10K bytes for the application program.

PRICING: IMS-2, IMS/VS, DL/1 DOS/VS, and DL/1 Entry and their supporting programs and features are standard IBM Program Products with full, centralized IBM Class A programming support. VANDL-1 is an IBM programming RPQ with Class C support.

Program	Number	Monthly License
IMS/VS	5740-XX2	$770
IMS/VS DC Feature	6001 or 6002	935
IQF for IMS/VS (DC)	6068,6069, or 6070	336
IMS-2	5734-XX6	616
IMS-2 DC Feature	6022-6024	700
IQF for IMS-2	6068, 6069, or 6070	336
GIS-2	5734-XX1	496*
GIS/VS	5740-XX7	864*
DL/1 DOS/VS	5746-XX1	330
DL/1 Entry	5746-XX7	300
VANDL-1	5799-AEY	378

*Numerous optional features can raise the total cost of GIS-2 and GIS/VS to more than $1500 per month.

INITIAL DELIVERY: IMS and the related products have been delivered in several releases, as follows:

Release Number	Distinguishing Characteristics	Date
2.0	Basic IMS-2	March 1971
2.1	3330 support	November 1971
2.2	Improved data communication service	August 1972
2.3	3270, virtual storage support	November 1972
2.4	Maintenance release	July 1974
1.0	IMS/VS for OS/VS1 and OS/VS2 Release 1.6	February 1974
1.1	VTAM support; enhancements to message format service and systems definition	April 1975
1.1.1	Service update	July 1975
1.0	DL/1 DOS/VS	December 1973
1.1	Support for variable-length segments; direct logging to CICS log tape; user exits for encoding	November 1974
1.0	DL/1 Entry	March 1975
1.0	VANDL-1	November 1972
1.1	CICS interface	January 1973
1.2	Maintenance release	–
1.3	3340 support	May 1974
1.4	Maintenance release	October 1974

CURRENT USERS: IBM will not release statistics on current numbers of users. However, Datapro estimates these figures to be approximately 450 for IMS/VS and IMS-2, and 150 for DL/1.■

IMS
IBM Corporation

The availability of DL/1 DOS/VS and DL/1 Entry seems to have encouraged a number of users of smaller systems, who could not conceivably have employed the full IMS, to stay with the IBM data base management philosophy. Many of the DL/1 users we spoke with said that if IBM had not introduced the smaller DBMS versions, they would have been forced to seek out suitable systems from an independent vendor. Although several of these users were still not happy about the amount of resources required for the smaller systems, they indicated that the facilities offered with the systems and the support they can expect from IBM formed the basis for their selections. They also noted that they feel they can grow with the smaller systems for a long time before it is necessary to think of changing. Again, these users emphasized that it pays handsomely to have experienced personnel on board. In some cases the users indicated that the training offered by IBM was adequate for their people and that the Data Base Administrator function was being handled by one of their own people whom they designated to be trained by IBM.

In all cases, the users we spoke to agreed that the installation of a DBMS should be accomplished slowly and, if possible, in stages, so that the true capabilities of the system can be understood before the user tries to move on to the more sophisticated features it offers. We think this is sound advice for any DBMS installation.

You're not likely to hear many comments from users about technical incomprehensiveness or inaccuracy of IMS. It's truly impressive in both concept and execution. But, a word of caution: don't underestimate, as many of the users we talked to did, the magnitude of the commitment necessary to implement IMS. The cost is considerable in terms of computer hardware, manpower, time, and applications development effort, not to mention the direct rental charge for the software itself. Keep in mind, though, that a major percentage of the cost of implementing IMS is a front-end expense for development work which will be roughly related to the complexity of your own applications.

Prospective data base management system users are considering a big step both in terms of the impact on their DP shops and the corporate implications to every user department. They owe it to their corporations to consider IMS (or DL/1) and to carefully weigh it against the other data base management systems available on the market. Many users will find that IMS is simply too large a dish to digest, or that it requires too much tender loving care, or that they cannot reasonably expect to get it up and running within their allotted time requirements. But others will find that it fits their installations just right. And for those users, IMS will produce a much more efficient data processing shop and will yield significant long-term benefits for the entire corporation.□

SYSTEM 2000
MRI Systems Corporation

MANAGEMENT SUMMARY

MRI's SYSTEM 2000 is a powerful and flexible data base management system that performs efficient retrievals and updates across the full range of data base applications. It operates on IBM System/360 and 370, Univac 1100 Series, and Control Data 6000, Cyber 70, and Cyber 170 Series computers. SYSTEM 2000's multiple indexing (a term utilized by MRI to describe what some other vendors call inversion and partial inversion) produces extremely fast retrievals and updates, even using complex selection criteria.

SYSTEM 2000 achieves rapid access to data base information by maintaining indices for data items, records, and structures. The data goes on to a serial file in compacted form, ready for retrieval. To process selection criteria, SYSTEM 2000 never searches through the data. Instead, it uses indices for rapid access to essential facts about the data records. An index of values for key data items supplies the locations of all records containing each value. An index of structural relationships enhances the selection capability by allowing users to request records having specified structural properties as well as the indicated values or combinations of values. The system also maintains passwords and authorities associated with each data item, guarding data integrity against unapproved modification, retrieval, or selection.

Increases in the speed of retrievals and modifications of the stored data can be achieved through thoughtful definition of the data base and skillful use of the supplied access languages. The inclusion of new records in the data base requires that the indices be updated also, but thereafter an index can be used to quickly locate the record without any additional overhead.

Creation of the SYSTEM 2000 data base entails describing the data using SYSTEM 2000's Data Definition Language (DDL) and then loading it through any of the system's data capture facilities. The system offers an overflow capability to accommodate exceptionally long data values. This enables the user to allocate space to an item based on the average length of the value instead of the longest value. The system will also group related data items into recurring records and organize the records into hierarchical structures, thereby reducing, in most cases, the volume of data passing through the system. The user has the ability to select only essential items for keying, and by anticipating future additions to the value index, he can fine-tune the everyday production environment with a minimum of complexity. SYSTEM 2000 allows queries with key or non-key selection criteria, and also allows items to change from non-key to key or vice versa without impacting user programs.

The data capture architecture of SYSTEM 2000 enables each installation to build its data base as the information ▷

SYSTEM 2000 is a generalized data base management system that features fast access either through its imbedded command language or through its interface to languages such as COBOL, FORTRAN, Assembler, and PL/1. SYSTEM 2000 operates in either an on-line or batch mode on IBM System/360 and 370, Univac 1100 Series, and CDC 6000, Cyber 70, and Cyber 170 computers.

CHARACTERISTICS

SUPPLIER: MRI Systems Corporation, 12575 Reasearch Boulevard, Austin, Texas 78766. Telephone (512) 258-5171.

BASIC FUNCTION: SYSTEM 2000 is a full-scale, generalized data base management system that supports multiple-indexed hierarchical data structures for efficient on-line and batch data processing. In all cases the user specifies which items are to be indexed. Numerous options add interface capabilities for COBOL, FORTRAN, PL/1, and Assembler languages in addition to the Self-Contained Query/Update language. The interface language is called Procedural Language Interface (not to be confused with PL/1) and allows the programmer to access files directly from a COBOL, FORTRAN, or PL/1 program. The system provides fast and easy access to data bases of all sizes and degrees of complexity.

OPERATION: SYSTEM 2000 consists of several features which can be used from remote or local terminals in either a batch or interactive environment. The user, after analyzing his information requirements, sets up a data base definition somewhat like the record formats used in layouts of conventional file management systems. The user also specifies structural relationships among records. A variety of loading facilities enter data into the data base–incrementally or all at once, with or without validation, in special streams or in existing formats, from a batch input device or a remote terminal.

The major features found in SYSTEM 2000 are:

- A Self-Contained Query/Update facility which permits immediate, direct updating of data base entries through an on-line keyboard or batch input.
- A Report Writer Feature which allows the user to define and generate as many as 100 reports in a variety of formats from a single pass of the data base indices under either batch or on-line processing.
- Procedural Language Interfaces for the use of COBOL, FORTRAN, PL/1, and Assembler which enable the user to access the data base through "get" and "put" macros.
- Audit facilities which preserve a machine-legible file of all update transactions and can be used with an archival copy of the data base for audit and backup.
- A multiple-thread feature which allows up to nine command streams to be handled simultaneously.
- Security provisions which offer password lockout of specific data items for selection, retrieval, or updating. ►

SYSTEM 2000
MRI Systems Corporation

and machine resources become available, therefore making lesser demands on the user who is just getting started in data base management, yet allowing him to grow within the system. The Host Language Interface offers an incremental load feature which handles additions to the data base from existing serial files. Optional validity checking features can be coded in the host language. In addition to the usual update commands, the Self-Contained Access Language includes two special commands, LOAD and UNLOAD, which read and write data values in an optimized loader string format. As soon as information enters the data base, it becomes available for immediate use. According to the vendor, a SYSTEM 2000 data base can actually reduce the physical storage space required for a body of raw data by as much as 40 percent. Some of this saved space is reallocated to use in the indices, but the vendor's claim is that the system still averages about a 15 percent saving in overall storage area.

English-like access languages enable users to express requests quickly and precisely. The Host Language Interface allows FORTRAN, COBOL, or PL/1 programmers to incorporate the organization and storage techniques of the SYSTEM 2000 data base into the procedural logic of a host program. The Self-Contained Query/Update Language offers execution modes designed for batched transactions. The SYSTEM 2000 Report Writer prepares up to 100 preformatted reports on a single pass through the data base. All the access languages share such significant features as: multi-threading, which serves requests in parallel, restricting other updates only during the application of an update transaction from one of several users simultaneously accessing a data base; a transaction log which assists the user in restoring a data base to its undamaged condition in case of system failure; and item and entry security at the entry level by data content.

USER REACTION

The 1975 Datapro survey of proprietary software users was conducted prior to MRI's release of Version 2.70 of SYSTEM 2000, and reflected some rather harsh user ratings, especially in the category of throughput. Since that time, the new version has been introduced and is currently in beta test at a customer site. The vendor reports that the initial load I/O time (for 20,000 logical entries) has improved by 73 percent over that of Release 2.60, and that the incremental load I/O time (for the same 20,000 entries) is 203 percent faster than under the previous release on this user's system. Loading of larger and more complex data bases is reported to have improved by over 300 percent. and retrieval by 200 to 300 percent in I/O efficiency. The vendor attributes these dramatic improvements to a new paging algorithm, new buffering techniques, and an extended set of tuning facilities incorporated into the new release.

In the 1975 Datapro survey, 13 users of SYSTEM 2000 responded. Their ratings of the package were mixed, with throughput/efficiency drawing a weighted average rating

- A "non-key where-clause" which allows users to include key or non-key items among their selection criteria.
- Tuning facilities which allow the data base administrator to improve the efficiency of his operation through operations such as modifying the buffer pool allocation and block size.
- A LINK feature in the Host Language Interface, which allows the user to establish temporary logical network relationships among records. Retrieval of one record can automatically initiate the retrieval of another–in another data base if required. Program logic within the system allows for the creation, modification, and deactivation of such associations, as well as opening and closing data bases as needed. This offers the user the ability to reorganize, restore, and access associated data bases independently.
- Teleprocessing interfaces to communications monitors such as MRI's TP 2000, IBM's CICS, and others, as well as with access methods such as TCAM. The TP 2000 interface supports the IBM 2740, 2741, and 3270 terminals and includes terminal display paging and input editing features.
- A Sequential File Feature which grants the user access to data bases residing on tape. Either the Report Writer or Host Language Interface programs can access such data bases.

The basic data storage technique used in the current SYSTEM 2000 version (Version 2.70) employs serial and multiple-indexed data structures. Indexing techniques facilitate high-speed retrieval and tend to conserve storage for redundant or missing items.

Variables which will influence the size of the data base include the lengths of data values within each record, the number of different data values in the data base compared to the total number of values, and the amount of indexing used.

The primary concept in the SYSTEM 2000 hierarchical file structure is the repeating group (analogous to a conventional record). Repeating groups nested within other repeating groups provide up to 32 levels of data structure with the possibility of unreleated records at each level.

HARDWARE/SOFTWARE REQUIREMENTS: On IBM computers, SYSTEM 2000 requires at least a System/360 Model 40 or System/370 Model 145 with a minimum of 256K bytes under OS, OS/VS, or VM. SYSTEM 2000 also operates under CMS. Sufficient on-line disk storage for the data base(s) and scratch files is required. Direct SYSTEM 2000 main storage requirements can vary from 140K to 200K bytes.

MRI recently announced that a DOS and DOS/VS version of SYSTEM 2000 will be available in the third quarter of 1976.

SYSTEM 2000 also runs under EXEC 8 on a Univac 1106, 1108, or 1110, using approximately 32K words of memory, and on a CDC 6000, Cyber 70, or Cyber 170 Series computer under SCOPE, KRONOS, or NOS, using approximately 20K words of memory.

PRICING: SYSTEM 2000 is available under two different arrangements. The External Use Agreement is intended for data services organizations that sell services externally. Initial payment is $25,000, covering the cost of installation, initial training, and availability fee. Maintenance service to the licensed organization is included for the life of the

SYSTEM 2000
MRI Systems Corporation

SYSTEM 2000 PRICES

Module	Paid-Up Lease	Permanent Monthly Lease Plan				Rental Plan	Annual Maintenance
		Number of Monthly Payments					
		12	18	24	36		
Basic SYSTEM 2000	$30,000	$2,665	$1,910	$1,550	$1,195	$1,200	$1,800
COBOL Procedural Language Interface	10,000	890	645	530	410	400	840
FORTRAN Procedural Language Interface	10,000	890	645	530	410	400	840
PL/1 Procedural Language Interface	10,000	890	645	530	410	400	840
Immediate Access	25,000	2,225	1,595	1,290	995	1,000	1,500
Report Writer	15,000	1,335	955	775	595	600	900
Sequential File	15,000	1,335	955	775	595	600	900
Accounting Log	2,500	220	165	140	110	100	300
Multi-Thread	20,000	1,775	1,310	1,085	870	800	2,400

of only 2.4, while overall satisfaction drew a 3.0. In contacting three additional users of the product, Datapro again noted that throughput/efficiency was the category which most users rated lower than the others, although the three users we interviewed were using earlier versions of the system.

Of the total of 16 respondents, one was using a Univac 1110, one a CDC 6600, and the remainder either an IBM System/360, Model 50 or larger, or a System/370, Model 155 or larger. The majority of the IBM users were using SYSTEM 2000 in a virtual environment, which could account for the large average memory usage of 292K bytes reported in the survey. The vendor claims that the average real memory requirement in a VS environment is only about 50K.

The 16 users had been using SYSTEM 2000 for an average of 17 months. The majority of the users praised the package for being flexible and for reducing programming time significantly. On the other hand, seven of the users criticized the product as being costly and using excessive resources.

MRI maintains that the problems identified by the users will be remedied with the formal release of version 2.70, due for release in early 1976. The MRI description of the improvements to the system with the new release is "dramatic." Many of the improvements implemented in the new release are attributed to the fact that MRI is "very responsive to customer needs," and the SYSTEM 2000 users we talked to seemed to agree.

Here's how the 13 survey respondents and the 3 users we interviewed by telephone collectively rated SYSTEM 2000:

	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	5	8	3	0	3.1
Throughput/efficiency	3	6	4	3	2.6
Ease of installation	5	7	3	0	3.1
Ease of use	1	12	2	0	2.9
Documentation	2	9	4	0	2.9
Vendor technical support	2	10	1	2	2.8
Training	1	7	4	0	2.8

*Weighted Average on a scale of 4.0 for Excellent.

As a group, data base management systems did not do as well as many of the other types of packages in the 1975 Datapro user survey. It appears that the users of these types of systems tend to demand more from the vendors inasmuch as the required investments both in money and manpower resources are so large. Thus, it is especially important for DBMS vendors to constantly address and correct the deficiencies in their products in the manner that MRI appears to be doing in the upcoming release of SYSTEM 2000.

Datapro recommends that anyone evaluating prospective data base management systems take the time to consider SYSTEM 2000 as a potential candidate that offers a number of noteworthy features. And, in the case of SYSTEM 2000, the recommendation applies to large-scale Control Data and Univac as well as IBM computer installations.□

contract. In addition, a royalty is charged and a minimum level of royalty per annum is required. Additional installation fees are applicable where versions for more than one hardware/executive combination are required.

The Internal Use Agreement provides an organization with SYSTEM 2000 for internal use only. Two lease plans and a rental plan are available at the prices shown in the table. The Paid-Up Lease allows use of the software for 50 years with a no-charge renewal option, and includes maintenance for the first year. The Permanent Lease Plan schedules payment over a period of 12, 18, 24, or 36 months; maintenance is included during the payment period.

The Rental Plan is a single month-to-month license agreement, cancellable with 30 days' notice; maintenance is included. The installation fee is $1,000 for the Paid-Up or Permanent Lease Plan and $2,500 for the Rental Plan. Multiple-installation discounts are available. GSA price schedules are available upon request.

INITIAL DELIVERY: Version 1 (fully indexed form), June 1970; Version 2 (partially indexed), July 1971. Version 2.7 (multiple-indexed) is scheduled for release in early 1976.

CURRENT USERS: Over 400 users, representing over 100 sites, throughout the United States, Canada, Europe, Australia, and Japan. Approximately 75 of these sites are within the United States.■

ADABAS
Software AG of North America

MANAGEMENT SUMMARY

ADABAS (Adaptive DAta BAse System) is an interesting and powerful data base management system that competes with IBM's IMS (Report 70E-491-01), Cincom's TOTAL (Report 70E-652-01), and several other systems (see the comparative charts in Report 70E-010-61). Pronounced "aid-a-base," the system was released in 1970 and became available in North America in 1972. Early in 1976 there were more than 130 users of ADABAS worldwide, including more than 75 in North America. The second International ADABAS Users' Conference, held in June 1975, was attended by more than 110 representatives of installations in the U.S., Canada, England, and West Germany.

While ADABAS ranks as a top-notch data base management system, few of its features offer unique capabilities (although one or two certainly *are* unique). Rather, the system has been put together with such thoroughness and skill that its most impressive features are not directly evident to the end user and appear in the form of reduced system overheads.

For example, ADABAS uses a data compression algorithm to load data into the data base. This simple technique is widely available on a stand-alone basis, but its incorporation into ADABAS as an integral function of the system typically results in about a 30 percent reduction in the total volume of the structured ADABAS data base as compared with the original size of the raw, unformatted data before being loaded. Typically, finished data bases in most other systems are larger than the input data used to create them, and the "expansion ratios" for such bases range from about 2:1 to as much as 8 or 10:1, depending upon the inherent complexity of the data relationships and the overhead imposed by the data base system itself.

Another unique feature of ADABAS is the separation of physical data storage from the representation of logical relationships in the data base. This concept permits the representation of completely general network structures and eliminates the necessity for data base reorganization—two characteristics not exhibited when hierarchical and chaining techniques are employed as in several other data base systems.

One of the most exciting potential benefits of ADABAS is related directly to the efficiency of the system and its logic design flexibility. While most of the parameters that define the system's capability are generally in the same ball park as those of other contemporary data base management systems, ADABAS has been designed to accommodate truly huge application environments (up to 4.2-billion-record data bases or larger). For example, ADABAS is nominally able to handle 255 files (or data bases). Each file can have a maximum of more than 16 ▷

ADABAS, a powerful data base management system, boasts a high level of craftsmanship in its implementation and wins praise from users for its flexibility and ease of installation. ADABAS runs on IBM System/360 or 370 computers under DOS, OS, DOS/VS, or OS/VS, and on Univac 9000 Series computers under the Disk Operating System.

CHARACTERISTICS

SUPPLIER: Software AG of North America, Inc., 1180 Sunrise Valley Drive, Reston, Virginia 22091. Telephone (703) 860-5050. German office: 61 Darmstadt, West Germany, Hilpertstrasse 20. British supplier: ADABAS Ltd., 3 Gray's Inn Square, London WCIR 5AH, England. Other offices in Brazil, Japan, South Africa, and Sweden.

BASIC FUNCTION: ADABAS is a data base management system with a number of utility programs for use under DOS, OS, or their VS counterparts, with BDAM used for data base generation and access. The system uses a variety of high-efficiency data management techniques and provides a generalized file-coupling capability. ADABAS operates as a "host" language system with high-level CALLS supported by ADAMINT, the ADABAS Macro Interface. A natural-language query capability, called ADASCRIPT, is supplied for "self-contained" operation. ADAWRITER, a report generator, is also provided. Data dictionary functions are provided by an interface to MSP's (Management Systems & Programming Limited, of London and New York) DATAMANAGER product. Interface modules also exist for communications monitors such as Intercomm, CICS, TSO, TASK/MASTER, ENVIRON/1, FASTER, and IMS/DC, and for report writers such as EASYTRIEVE, MARK IV, SCORE, Quikjob, and Data Analyzer.

OPERATION: ADABAS operation requires that the user perform the following steps (presented in the typical development sequence, although considerable overlapping normally takes place):

- The applications are defined in terms of functional requirements and types of data to be processed. Individual applications programs are then written in COBOL, PL/1, FORTRAN, or Assembler language to perform these applications, with calls inserted at appropriate points for the ADABAS access commands. Existing sequential processing programs using the data must be recompiled after some modification (consisting mostly of changing the Reads and Writes to ADABAS calls using the ADAMINT facilities).
- All of the individual data requirements for each application should be coordinated into an overall data base requirement. This task is most appropriately handled by an individual serving as the Data Base Administrator. This function is much more nearly administrative in nature for ADABAS than for systems such as IMS, where the "administrator" duties include complex data base design responsibilities as well.
- If communications or interactive query requirements are present, a choice must be made among the ▶

ADABAS
Software AG of North America

million records, and each record description can have 500 types of data fields, of which 200 can be key or "descriptor" fields. In addition to ADABAS' logical "network" capability, any number of geographically distributed data bases can be interfaced in support of informational requirements for combinations of data.

The basic file structure used by ADABAS is a full inversion of selected fields in a file which lets the user identify the key elements for retrieval, but is developed in a way that minimizes storage and processing overhead. Records in each file can be logically coupled to any number of records in up to 80 other files. These coupling relationships are defined after initial loading of the file, and the parent-child relationship is specified as part of the inquiry, thus resulting in a previously unattainable after-the-fact "network" capability.

The design of ADABAS also addresses the problem of ease of use and flexibility—common problems with sophisticated data base management systems. ADABAS files are defined in terms of fields and records, and five basic commands (load, modify, read, find, and delete) provide the full range of data management capabilities. The user can change keys for data retrieval, add fields to files, add files to the data base, and add relationships between files, all without impact on the user programs and without reloading the data base itself. In addition, a reasonably complete set of more than 200 diagnostic messages is included in the basic system to facilitate use of ADABAS and to aid in debugging.

Data base integrity and security are also addressed in the ADABAS design. Complete checkpoint and recovery procedures are built in to protect against program failure and machine failure. User security can be provided by passwords at up to 15 levels to protect against unauthorized accessing and updating at both the file and field levels.

ADABAS does not currently include a native data communications monitor for remote access. However, Software AG of North America has developed and maintains interfaces for CICS, TSO, Intercomm, TASK/MASTER, FASTER, ENVIRON/1, and IMS/DC. ADABAS can be used simultaneously by user systems in any number of batch and teleprocessing regions.

For stand-alone use, ADABAS provides a native-language query capability called ADASCRIPT, which can be used as a batch program or interactively under any communications monitor. A report generator called ADAWRITER is also provided with ADABAS, and interfaces have been built for many commercially available report writers.

What becomes clear as the ADABAS system is studied is that a truly remarkable implementation of many worthwhile data base concepts has been made in ADABAS. In fact, it seems that there are no poorly implemented aspects of ADABAS's basic design. The

independent communications monitors. ADABAS has interfaces for most of the commercially available monitors, as noted earlier in this report.

- At this point the actual data base can be put together, using special-purpose off-line utility programs. The existing data base or raw data is presented as standard sequential files to the system.

The ADABAS data base is stored on disk using IBM's standard Basic Direct Access Method (BDAM) and consists of three main components: 1) the actual data stored on disk with a unique Internal Sequence Number (ISN) assigned to each record, 2) an "Associator" consisting of indices and structures defining the logical relationships between files in the data base, and 3) an Address Converter that translates the logical addresses in the Associator into physical addresses of records.

The data storage portion of ADABAS permits up to 4.1 billion records in as many as 255 logical files to be treated as an entity. Each file can contain a maximum of 16.8 million records, and each file can hold 500 distinct field names, 200 of which can be designated as descriptors. The maximum individual value length is 253 bytes. Each record can have fixed-length or variable-length fields containing alphanumeric, binary, fixed-point, or decimal packed/unpacked data.

The data base structure is in the Associator, which holds the indices and coupling relations used for record retrieval. Sorted ISN lists are maintained for user-specified key fields, known as "descriptors," which are used in retrieval requests and for coupling files together.

The Associator information is updated automatically, and this process is approximately linear. Each ISN is 3 bytes (24 bits) long and is used to reference up to 16,777,214 records in each file. As the data base is built (or modified), a directory is maintained and can be read by the user. A "histogram" of descriptor values can be obtained at any time without the need to read data records.

When a retrieval request is processed, the user-specified search criteria are matched against the appropriate inverted fields, and the most economical search strategy is chosen. The shortest list of qualifying ISN's is retrieved first, thus minimizing the number of disk accesses required to satisfy the request. Software AG's statistics indicate that, because of multi-level indices and buffer management techniques, a simple query (single criterion) will rarely require more than two accesses. The worst case for reading data records will average two accesses per record (Address Converter and Data Storage). Sets of retrieved records can be processed in sorted order on up to three keys.

ADABAS is supported by several utility subsystems. These include:

- Loader utilities to build and add to the data base using complex input data edit criteria and a data compression technique. The data compression routines eliminate the storage of leading zeros for numeric fields, trailing blanks for alphanumeric fields, and imbedded null fields. All of these capabilities also exist in the core-resident non-overlaid system for on-line updating. In ordinary business environments, the ADABAS data compression results in about a 50 percent reduction in the volume of stored data as compared with the total length of raw input data. (Indices and related information add an overhead that can restore the final data base size to about the full size of the unedited input data; typical final size,

ADABAS
Software AG of North America

fundamental weakness of classic hierarchical, semi-hierarchical, or network file structures (such as that of ADABAS) stems from the fact that a trade-off must be made between efficiency of data organization for retrieval purposes and the complexity of creating and updating the data base. Through clever programming techniques, ADABAS has been able to minimize the data structure overhead and consequently reduce the complexity of modifying the data base, allowing efficient retrieval without undue update overhead.

Software AG has put more than 50 man-years of development effort into ADABAS to date, and a considerable number of users are finding to their satisfaction that the resulting system has considerable merit. With more than 75 North American users to date, ADABAS seems to have gotten on base despite the two strikes it initially had against it: its high price tag and the not-invented-here (U.S.) syndrome. If a medium-to-large scale data base management system is one of your needs, then an in-house demonstration of ADABAS on your own data is well worth having.

A comparison of ADABAS with other current data base management systems can be found in Report 70E-010-61.

USER REACTION

In the 1975 Datapro survey of proprietary software users, nine respondents rated ADABAS on the basis of their experience with it. The average length of time in use for these respondents was 14 months, and the average main memory usage attributed to ADABAS was 194K bytes. Here's how these nine users rated the system:

	Excellent	Good	Fair	Poor	WA*
Overall satisfaction	5	3	1	0	3.4
Throughput/efficiency	2	4	1	0	3.1
Ease of installation	7	1	0	0	3.9
Ease of use	3	4	1	0	3.3
Documentation	1	2	5	0	2.5
Vendor technical support	2	4	2	0	3.0
Training	4	1	2	0	3.3

*Weighted Average on a scale of 4.0 for Excellent.

As you can see, ADABAS was rated very highly by these respondents in all areas except the category of documentation. It is a well-known fact that the documentation for a data base management system takes a great deal of time and effort in order to satisfy both the sophisticated and the novice user of the system. In this respect, the ADABAS documentation was rated fairly low by both types of users. It is only fair to say that as of December 1975, Software AG of North America has taken steps to update the documentation with Version 3.2 of the system. Datapro reviewed this new documentation, and it is obvious that more concern has been given to explanations, examples, and diagrams depicting the logic involved in step-by-step processes.

however, is about 70 percent of that of the raw input data).

- A Couple Utility to logically connect a pair of files. This is a two-way connection, providing the file networking capability in ADABAS; either file can be used as the "parent" and the other as the "child" in any retrieval. A record can be coupled to any number of records in another file and any file can be coupled to as many as 80 other files.

- A Restart Utility to resume ADABAS operations following an interruption, including resetting the data base and Associator to the most recent checkpoint for rerun. ADABAS can automatically regenerate lost data, rerun an application program from a previously defined checkpoint, and recover from machine failures.

- Modification Utilities to establish a field as a descriptor, release a descriptor, release coupling, define new fields, reorder data base storage, and add new storage devices.

- A Report Utility to provide information about data base definition and space utilization.

Two unique aspects of ADABAS are 1) a phoneticized retrieval capability, and 2) a data encryption or cypher option for ensuring data base security. Both of these functions are noteworthy, but they will appeal only to a subset of general-purpose data base users. ADABAS provides up to 15 levels of security for the protection of data from unauthorized accessing and updating.

Two additional aspects of ADABAS performance must be considered: 1) data base creation or modification, and 2) ADABAS operation for processing retrievals.

For data base creation, Software AG states that ADABAS has been timed to compress, invert, and load from 200,000 to 400,000 records per CPU hour on a System/370 Model 155.

For retrieval, many parameters are involved; but as an example, a 360/50 processing a 1-million-record data base with 200 fields per record against a 5-element set of search criteria took 4 seconds to determine the number of qualifying records, retrieve their ISN's, and read the first record. The second and subsequent records were returned within the average access time of the secondary storage device, as the list of qualifying ISN's resides in main memory after the initial ADABAS analysis.

HARDWARE/SOFTWARE REQUIREMENTS: ADABAS consists of a 130K-byte non-overlaid nucleus with about a 1-million-byte disk-resident library including the utilities. The nucleus plus a variable-sized buffer area (minimum 30K) must be resident in main memory at all times. In a multi-user environment, a re-entrant 8K to 16K front-end module is used to queue access requests, while accesses are handled in a single-thread or non-reentrant mode. ADABAS can be run in a 170K partition or region under DOS, OS, or their VS counterparts on an IBM System/360 or 370, or on an IBM-like Univac 9000 Series or Siemens 4004 system under the Disk Operating System (PBS for Siemens). IBM 3340, 2314, 2305 or 3330-type disk drives can be used for the data, Associator, and work files, and one magnetic tape unit is required for data protection.

The size of the data base itself depends upon the amount of raw data and the degree to which it is compressible (i.e., number of null fields, etc.). Typically, the data can

ADABAS
Software AG of North America

The rating in the overall satisfaction category was only 0.1 point short of attaining the required 3.5 level that is a prime requirement for Honor Roll status. Of course, the 2.5 rating in documentation would have still been a roadblock to achieving this goal, but the 3.4 rating in overall satisfaction definitely shows that the users of the package have a very high regard for its performance. This conclusion is strengthened by the fact that all the other categories, besides documentation, were rated 3.0 or above.

As in past Datapro surveys, the strongest point indicated by the ADABAS users was the system's ease of installation, where seven of the nine respondents rated the package as excellent and one as good, while the ninth did not rate this category. Moreover, five of the nine respondents indicated that the product was installed and performed immediately as advertised, and two others reported that the system eventually performed as advertised.

Other strong points noted by the respondents to the survey were the flexibility of the product and the savings in programming time in developing the DBMS environment. Two of the users, however, felt the product was costly.

Judging from the very positive reactions of its users, ADABAS should certainly be included in any serious evaluation of data base management systems. □

be compressed to about 50 percent of its raw input size for commercial applications. In addition to this, an Associator is built containing the data base structure relationships. The Associator normally requires three to four bytes per descriptor occurrence. The final data base size tends to vary from about 50 percent to 100 percent of the size of the raw input data.

PRICING: Software AG of North America, Inc. provides ADABAS on a permanent-license basis for $120,000 for the first CPU, including all normal training. The second CPU license costs $60,000. Lease plans range from month-to-month at $5,000 per month, to five years at $2,500 per month, with a $2,000 training fee. Purchase options on all leases allow a certain percentage (50 to 65) of the accumulated payments to be applied to purchase. Future releases (expected to be about one per year), along with maintenance and enhancements, are included in the cost of the lease and are provided for $6,000 per year with purchase. ADAMINT is available for a $6,000 one-time charge and has a $600 annual maintenance fee.

A user-paid demonstration is available for $2,500. Software AG will convert part of the user's data base to ADABAS, run one or two of the user's existing sequential processing programs against it at the users's own site, and demonstrate the use of ADABAS commands for retrievals from that file.

INITIAL INSTALLATION: Version 1 (Basic ADABAS)–March 1971; Version 2–September 1972; Version 3–January 1974.

CURRENT USERS: Over 130 installations, with more than 75 in North America. ■

First Class Mail
Permit No 178
Delran, N J 08075

Business Reply No Postage Stamp Necessary if Mailed in the United States

Postage will be paid by:

datapro

DATAPRO RESEARCH CORPORATION
1805 Underwood Boulevard
Delran, N J 08075

Save time and money with these datapro $12 reports

Revised and expanded reports that include many new systems and models for 1977

53 objective reports on EDP and office products to help you make better planning and buying decisions.

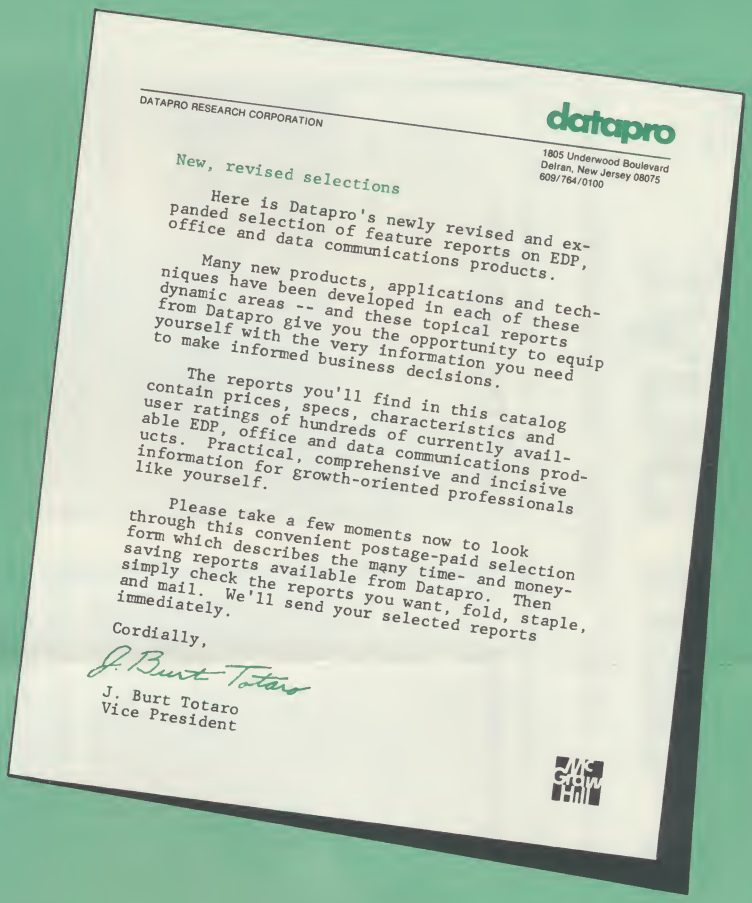

DATAPRO RESEARCH CORPORATION

datapro

1805 Underwood Boulevard
Delran, New Jersey 08075
609/764/0100

New, revised selections

Here is Datapro's newly revised and expanded selection of feature reports on EDP, office and data communications products.

Many new products, applications and techniques have been developed in each of these dynamic areas -- and these topical reports from Datapro give you the opportunity to equip yourself with the very information you need to make informed business decisions.

The reports you'll find in this catalog contain prices, specs, characteristics and user ratings of hundreds of currently available EDP, office and data communications products. Practical, comprehensive and incisive information for growth-oriented professionals like yourself.

Please take a few moments now to look through this convenient postage-paid selection form which describes the many time- and money-saving reports available from Datapro. Then simply check the reports you want, fold, staple, and mail. We'll send your selected reports immediately.

Cordially,

J. Burt Totaro

J. Burt Totaro
Vice President

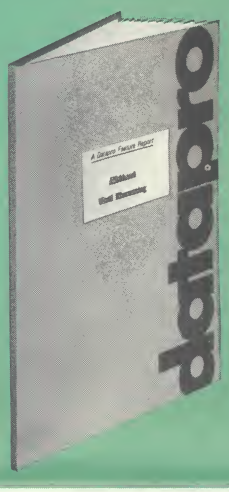

- Determine product performance, without buying, via published users' ratings.
- Identify best price-performing products by direct comparison of specs and costs.
- Control costs by improving the fit of products to your needs.
- Eliminate costly mistakes resulting from improper product selection.
- Outpace competition by increasing awareness of industry trends.
- Identify new or growing markets.
- Cost-optimize EDP functions.
- Save time screening vendors.

SEE INSIDE
for complete descriptions of all 53 reports

Reminder Before Mailing

After selecting reports, make sure that your name and address are complete on the lower panel on the back page, and the entire postage-paid form is folded and stapled with this side out.

First Class Mail
Permit No. 178
Delran, NJ 08075

Business Reply No Postage Stamp Necessary if Mailed in the United States

datapro

DATAPRO RESEARCH CORPORATION
1805 Underwood Boulevard
Delran, New Jersey 08075

SEAL THIS EDGE
WHEN FOLDED

more time and money-saving reports inside!

☐ All About Telephone Answering & Diverting Devices

This 29-page report surveys the applications and features of telephone answering and call diverting devices marketed by 20 manufacturers or suppliers. Accompanying charts list the characteristics of over 70 announce-only, announce and record, and call diverting units sold over the counter or through a dealer. *Selected from Datapro Reports on Office Systems.* $12.

☐ All About Facsimile Equipment

The increasing demand for "instant" communications, coupled with a decline in the speed of mail service and higher voice telephone costs should have a beneficial effect on the sales of facsimile equipment. Market forecasts place such sales at the $250 million level by the end of the decade. This 27-page report compares and rates the features and characteristics of 70 models from 15 facsimile equipment suppliers. *Selected from Datapro Reports on Office Systems.* $12.

☐ How to Select a Telephone System

Deciding on a specific new telephone system may be one of the most important systems and equipment decisions you'll be asked to make in the near future. This 22-page landmark report can save you weeks and months of time in the work leading up to system selection by providing you with a master plan, including charts, forms, and check-lists. *Selected from Datapro Reports on Office Systems.* $12.

EDP peripherals

These 9 up-to-date reports can help you optimize the cost of structuring your data processing capabilities. You'll find detailed information about performance, pricing, installation, and maintenance.

☐ All About Magnetic Tape and Disk Pack Maintenance Equipment

A well-planned media maintenance program can lead to significant cost savings and increased operational efficiency. This 10-page report describes the currently available equipment for cleaning, inspecting, testing, and certifying disk packs and half-inch magnetic tape. Users' ratings and detailed prices are included wherever possible. *Selected from Datapro 70.* $12.

☐ All About Data Collection Equipment

Data collection is the catch-all phrase for data input devices normally used in environments other than computer data preparation areas. The equipment ranges from central recording systems for reporting attendance and work-in-progress information, to hand-held tape recorders for inventory taking. This 20-page report surveys the characteristics of 45 portable and stand-alone devices, and data collection systems available from 31 vendors. *Selected from Datapro 70.* $12.

☐ All About Plug-Compatible Disk Drives

Which, if any, of the numerous IBM-compatible disk subsystems represents the best buy for your installation? Or, is there a plug-compatible drive that can be used with your Univac, Honeywell, or Xerox computer system? This 23-page report helps you find the answer with practical guidelines, specifications of 59 disk subsystems from 14 vendors, and users' ratings of more than 2000 drives. *Selected from Datapro 70.* $12.

☐ How to Select and Use Add-On Memory

Add-on main memory from independent vendors can substantially improve the performance and cost-effectiveness of most computers. This report will help you evaluate the pros and cons of add-on memory, benefit from the experience of hundreds of users who reported on over a hundred million bytes of independent memory, and select the unit that's best for you. *Selected from Datapro 70.* $12.

☐ All About Graphic Display Devices

Available in a wide range of capabilities, graphic display devices provide the most powerful interactive facilities of any class of man/computer device. This 20-page report presents the pertinent characteristics of 53 devices from 18 vendors and includes a survey of the experience of 40 users. *Selected from Datapro 70.* $12.

☐ All About Optical Readers

This 36-page report discusses the position of optical readers in the world of data processing today and offers straightforward advice to help you determine whether this means of data entry can be used effectively in your organization. Included are tabulations of the experience of 131 optical reader users, and detailed characteristics of 101 currrent optical character, mark, and bar code readers from 41 manufacturers. *Selected from Datapro 70.* $12.

☐ All About Voice Response

This report forms a comprehensive overview of voice response techniques and equipment, including an analysis of user experience representing over 132,000 calls per day, and characteristics and pricing on over 21 systems and components. *Selected from Datapro 70.* $12.

☐ How to Select and Use Data Entry Devices

In view of the variety of data entry devices now on the market, "data entry bottlenecks" needn't be as much of a problem as it once was. This 16-page report provides a comprehensive look at the various types of data entry devices available, and how to select and use them effectively. *Selected from Datapro 70.* $12.

☐ All About Computer Output Microfilm

This 24-page report describes the characteristics, applications, advantages, and disadvantages of COM, and presents concrete equipment selection guidelines plus the results of a survey of the experience of 135 COM users. Comparison charts summarize the capabilities of 29 commercially available COM recorders from 12 vendors. Descriptions of microfilm formats and film types are included. *Selected from Datapro 70.* $12.

other important datapro reports

☐ A Buyer's Guide to Electronic Funds Transfer Terminals

This 24-page buyer's guide presents some general guidelines on what to look for when buying manned remote terminals. Also, you'll find detailed product reports describing the features and characteristics of currently available equipment. *Selected from Datapro Reports on Banking Automation.* $12.

☐ All About Integrated POS Systems

This report describes major integrated point of sale systems currently available for the general retail merchandising and food industry. It provides both a general understanding of these systems and detailed information about specifications and capabilities of each product. *Selected from Datapro Reports on Retail Automation.* $12.

☐ A Buyers' Guide to Electronic Cash Registers

Electronic cash registers, the most direct replacements for conventional electromechanical registers, represent all the automation that's needed in many retail organizations. This 24-page buyer's guide, which includes detailed comparison charts, will fill you in on what's available and how it meshes with your requirements. *Selected from Datapro Reports on Retail Automation.* $12.

☐ All About Photocomposition

Once the exclusive domain of the commercial or newspaper printer, photocomposition is now entering the ofice environment via a number of relatively low-cost and easy-to-operate models. This 27-page report reviews 66 models of photocomposers from 18 manufacturers, ranging from smaller, keyboard-operated models to the very large, computer-controlled or tape-fed systems. *Selected from Datapro Reports on Office Systems.* $12.

☐ All About Micrographics

This 55-page report summarizes the characteristics and features of more than 200 microform readers and reader/printer from 49 vendors, including hand-held, portable, desk, console, and automatic models. If you are considering the purchase of a microform system, or if you need to keep current about significant improvements in product design, this report is an essential aid. *Selected from Datapro Reports on Office Systems.* $12.

To order: see inside

If you subscribe to more than one Datapro information service, or in certain other cases, you may likely receive a duplicate of this announcement. If so, we hope you will pass it along to an interested associate.

datapro

DATAPRO RESEARCH CORPORATION
1805 Underwood Boulevard
Delran, NJ 08075 Tel. 609-764-0100

BULK RATE
U.S POSTAGE
PAID
PHILADELPHIA, PA 19140
Permit No. 1325

NOTE: Please fill in name and address below. If address label is already affixed, please make any changes or corrections below.

Name _____

Title _____

Company _____

Address _____

City _____

State _____ Zip _____

Phone _____

Check what $12 can buy!

unbiased EDP and office product reports: characteristics, prices, users' ratings, comparisons, more.

minicomputers

Whether your interest is in distributed processing, or microprocessors, minis and small business computers per se, you'll find plenty of money-saving information about this fast-changing industry segment in these comprehensive reports.

☐ **All About Small Business Computers**

This comprehensive report is designed to help you select and apply low-cost business data processing systems. The characteristics and prices of over 200 current systems from over 96 vendors are reported in detailed comparison charts, and the experience of hundreds of small business computer users is analyzed. *Selected from Datapro 70.* $12.

☐ **All About Minicomputers**

This comprehensive report is designed to aid you in understanding the rapidly proliferating minicomputers and selecting the one that can best satisfy your requirements. You'll find detailed comparison charts covering the characteristics of over 200 current minicomputers from over 69 manufacturers, as well as users' experiences with thousands of minicomputers. *Selected from Datapro 70.* $12.

☐ **All About Microcomputers**

This 32-page report features detailed comparison charts that describe characteristics and pricing of current microcomputers from 67 vendors, plus new information on development systems and kits. *Selected from Datapro Reports on Minicomputers.* $12.

☐ **All About Microprocessors**

This completely revised, 18-page report details characteristics of current microprocessors from 27 manufacturers. You'll also find straightforward guidelines to help you understand and apply these compact, low-cost CPU's. *Selected from Datapro Reports on Minicomputers.* $12.

☐ **All About Minicomputer Disk Storage**

This 54-page report provides minicomputer users and suppliers with a comprehensive survey of 146 disk storage units. The comparison charts describe the current product lines of 47 suppliers, and the accompanying text offers the reader an enlightening guide for evaluating disk and drum storage performance. *Selected from Datapro Reports on Minicomputers.* $12.

☐ **All About Minicomputer Floppy Disk Storage**

Floppy disk drives have emerged as a distinct type of peripheral device, with an identity and set of performance parameters entirely different from those of cartridge or pack drives. This 22-page report examines the floppy disk and puts its performance into perspective. 78 products from 42 vendors are detailed in accompanying comparison charts. *Selected from Datapro Reports on Minicomputers.* $12.

data communications

These latest evaluations of data communications equipment can help you avoid costly mistakes in planning, installing or updating a system. These 11 valuable reports provide the facts you need to make cost-effective planning and buying decisions.

☐ **All About Remote Batch Terminals**

The path to high-volume transmission and reception of data at locations remote from your computer center is a confusing one. This 35-page comprehensive report provides detailed specifications of 109 remote batch terminals from 43 vendors, presents the results of an extensive user survey covering more than 1,700 installed terminals, and offers practical guidelines for selecting a batch terminal that will meet your specific needs. *Selected from Datapro 70.* $12.

☐ **All About Data Communications Facilities**

This unique 26-page report summarizes the communications facilities and rate structures of AT&T, Western Union, and the currently active microwave and satellite carriers. You'll also find the results of a survey of user experience with about 10,000 communications links, plus guidance in selecting the facilities that will most effectively satisfy your needs. *Selected from Datapro 70.* $12.

☐ **All About Alphanumeric Display Terminals**

This comprehensive 54-page report provides detailed specifications of 182 alphanumeric display terminals from 73 vendors, presents the results of an extensive user survey covering more than 14,000 installed terminals, and offers practical guidelines for selecting a display terminal that will meet your specific needs. *Selected from Datapro 70.* $12.

☐ **All About Teleprinter Terminals**

This 42-page report provides detailed specifications of 149 teleprinter-style communications terminals from 57 vendors, presents the results of an extensive user survey covering more than 11,000 installed terminals, and provides guidance in selecting the most suitable terminals for your needs. *Selected from Datapro 70.* $12.

☐ **All About Remote Computing Services**

This comprehensive 41-page report explains both interactive time-sharing and remote batch processing, discusses their advantages and drawbacks, summarizes the services offered by 100 remote computing companies, tells how users rate the leading companies, and provides guidelines for selecting a suitable supplier and getting started. *Selected from Datapro 70.* $12.

word processing

Activity in the word processing area continues at a frantic pace, with new product announcements appearing at a rate of one per month. Some 50-odd companies are now actively competing in the market—the overall business climate resembling that of computers in their heyday period of the early sixties. In these reports, you'll find specifications, features, costs, advantages and limitations of hundreds of alternative word processing products. Presented in concise text and extensive comparison charts.

☐ **All About Automatic Typing and Text Editing Systems**

This comprehensive 50-page survey summarizes the characteristics of over 100 models of automatic typing and text editing systems from 60 vendors. Detailed comparison charts describe each model's basic configuration, pricing, storage medium and capacity, editing features, printer characteristics and much more. *Selected from Datapro Reports on Office Systems.* $12.

☐ **All About Dictation Equipment & Systems**

The applications and use of dictation equipment in today's WP environment are discussed in this over 45-page Datapro report, and the features of over 100 central, desktop and portable dictation units are tabulated. *Selected from Datapro Reports on Office Systems.* $12.

☐ **All About Office Electric Typewriters**

Although automatic typewriters and text editors are making inroads into the office environment, electric typewriters are still the backbone of any office operation. New entries in the single element, Selectric-type segment of the $500-million-per-year market make it a good time to review your present typing equipment needs. Thirty-three models from 11 manufacturers are examined and compared. *Selected from Datapro Reports on Office Systems.* $12.

Analyses of Selected Word Processing Systems

You can save hundreds of hours of investigation with these objective, in-depth word processing system analyses. Each report describes the different systems in detail, including

office copiers

Copier usage continues to spiral, and with it, copying costs. These unbiased reports will give you the facts and background you need to identify and select the most cost-effective equipment for your needs.

☐ **All About 150 Office Copiers**

Here's a condensed one-for-one look at over 150 office copier models offered by 30 vendors that make up today's proliferating office copier market. In 50 pages you'll find concise text and extensive comparison charts to explain what's available today, and how it performs. *Selected from Datapro Reports on Office Systems.* $12.

Analyses of Selected Office Copiers

These three money-saving reports provide detailed analyses of major entries in the plain paper copier market—logically grouped to help you understand your options in the low, intermediate, and high volume areas. Specifications, price-performance measurements, operating characteristics, and much more.

☐ **Selected Office Copiers: Low Volume**

Detailed analyses of five major copiers intended for users needing 5,000 to 20,000 copies per month. All priced under $5000. Included are: Cannon NP70, Saxon 1, Sharp SF-720, Toshiba BD-702A and AB Dick Bond 901. 15 pages. *Selected from Datapro Reports on Office Systems.* $12.

☐ **Selected Office Copiers: Intermediate Volume**

Detailed analyses of three major copiers selling under $25,000. Suited for users needing 20,000 to 40,000 copies per month. Included are: IBM Copier II. Savin 750 and Xerox 3100. 15 pages. *Selected from Datapro Reports on Office Systems.* $12.

☐ **Selected Office Copiers: High Volume**

Detailed analyses of three major copiers for those needing 40,000 to 60,000 copies per month. Prices start at about $25,000. Included are: Eastman Kodak Ektaprint Series, IBM Series III, and Xerox 3600. 26 pages. *Selected from Datapro Reports on Office Systems.* $12.

general-purpose computers

☐ **All About Minicomputer Add-On Memories**

Whether the memory product "adds on" via a cable-connected box or "adds in" by inserting circuit cards in your minicomputer chassis, independent memory vendors can help you expand your system at reasonable cost—if you have one of the more popular minicomputers. This 26-page report surveys 107 memory products from 20 vendors. *Selected from Datapro Reports on Minicomputers.* $12.

☐ **All About Minicomputer Printers**

Until everyone can walk around with a vest-pocket display link to a computer data base, the printer remains the primary means for retaining information retrieved or generated by the computer. This 58-page survey looks at the products of 87 vendors in the form of comparison charts that detail the characteristics of 252 printers or printer families. It also reflects the progress that printer suppliers have made in solving the continuing problem of delivering increasing performance at decreasing cost. *Selected from Datapro Reports on Minicomputers.* $12.

software

In these reports, Datapro can help you determine the performance of proprietary software packages without buying, via published ratings from the largest user survey ever conducted. Make these two popular reports your fingertip guide to software selection.

☐ **User Ratings of Proprietary Software**

This 22-page report contains the results of Datapro's fourth annual survey of proprietary software users, conducted in cooperation with *Datamation.* You'll find extensive tables showing how thousands of users rated 250 individual packages, plus financial insights and the 1976 Software Honor Roll. *Selected from Datapro 70.* $12.

☐ **A Buyer's Guide to Data Base Management Systems**

This 32-page report presents useful guidelines to aid you in preparing, evaluating and selecting a data base management system. It compares the key characteristics of major systems and provides an in-depth evaluation of these systems: ADABAS, IDMS, IMS, SYSTEM 2000, and TOTAL. *Selected from Datapro 70.* $12.

☐ **All About Modems**

Data communications users can select from an overwhelming variety of modems offered by dozens of independent vendors and by the telephone companies. This in-depth, 48-page report summarizes the characteristics of more than 300 current modems from 43 suppliers, provides useful selection guidance, and includes a unique survey of users' experience with over 15,500 modems. *Selected from Datapro 70.* $12.

☐ **All About Data Communications Multiplexors**

This 25-page report explains the techniques and applications of multiplexing, summarizes the characteristics of 67 current units from 24 vendors, analyzes the experience of multiplexor users, and provides useful selection guidance. *Selected from Datapro 70.* $12.

☐ **A Buyer's Guide to Data Communications Monitors**

This 30-page report contains straightforward guidelines to aid you in evaluating and selecting a data communications monitor, compares the key characteristics of leading systems and provides an in-depth evaluation of these systems: CICS (all five versions), ENVIRON/1, Intercomm, TASK/MASTER, and the Teleprocessing Interface System (WESTI). *Selected from Datapro 70.* $12

☐ **All About Communications Processors**

This 32-page report analyzes the benefits and drawbacks of the programmed communications processor concept in general; it also describes the front-end environment in detail, reports the experience of 256 users, and in detailed comparison charts summarizes the capabilities of 84 communications processors. *Selected from Datapro 70.* $12.

☐ **How to Analyze Your Data Communications Needs**

This 10-page report explains what data communications can do and how you can put it to work in your organization. You'll find straightforward procedures for determining your communications needs and implementing an effective system to satisfy them. *Selected from Datapro 70.* $12.

☐ **IBM Synchronous Data Link Control (SDLC)**

A six-page presentation of IBM's SDLC line protocol, a de facto standard for data communications. *Selected from Datapro Reports on Data Communications.* $12.

individual management summaries, operating codes and characteristics, specifications, photos, purchase/lease pricing and more. Logically grouped to help you identify and compare the equipment best suited to your needs.

☐ **Word Processing: Selected Stand-Alone Printing Systems.**

Here's an in-depth examination of four contenders in this popular WP area: IBM Magnetic Card Typewriter line, Redactron Magnetic Tape Editing Typewriter, Redactron Magnetic Card Editing Typewriter, and Xerox 800 ETS (Electronic Typing System). 23 pages. *Selected from Datapro Reports on Office Systems.* $12.

☐ **Word Processing: Selected Stand-Alone Video Printing Systems.**

Added flexibility is inherent in these three video-based product lines. Detailed analyses of: Lexitron Videotype Text Processors, Vydec 1146, 1200, and 1400, and the Wang Laboratories Word Processor. 19 pages. *Selected from Datapro Reports on Office Systems.* $12.

☐ **Word Processing: Selected Shared Logic Multi-Terminal Systems.**

Large WP users will find a great deal of power and sophistication in these three minicomputer-based systems. Analyzed in-depth are: Daconics Visual Type 5400, ICS Astrocomp, and LCS Compu-Text. 18 pages. *Selected from Datapro Reports on Office Systems.* $12.

telephone/voice communications/ facsimile Specifications, features, costs and more:

☐ **All About Intercom Systems**

This 38-page report covers the features and characteristics of over 100 intercom systems designed for office or business environments, supplied by 25 different vendors. Such dedicated systems can increase overall efficiencies by decreasing employee walks to discussion sites and lessening telephone call-back charges, as well as improving information flow within the company. *Selected from Datapro Reports on Office Systems.* $12.

(continued on other side)

These important reports will keep you up to date, and help you select best-suited EDP equipment, media, and supplies. You'll find hundreds of users' ratings and in-depth product and vendor analyses.

☐ **Directory of Suppliers**

This 174-page directory contains informative profiles of 850 companies that offer EDP products and services of all types, plus guidelines for interpreting the information and selecting qualified suppliers. You'll find entries describing each company's location, size, executives, financial status, ownership, sales and service organization, and product line. *Selected from Datapro 70.* Special price: $15.

☐ **User Ratings of General-Purpose Computing Systems**

This eight-page report conveys the results of Datapro's 1976 survey of general-purpose computer users. Extensive tables summarize the experience of over 2000 computer system users. The users' ratings pinpoint the strengths and weaknesses of each mainframe manufacturer's equipment, software, and support. *Selected from Datapro 70.* $12.

☐ **IBM System 370**

This 71-page report describes and analyzes all of the System/370 equipment and software, including such recent developments as the 115-2, 125-2, 158-3 and 168-3, and summarizes the results of a survey of 863 System/370 users. *Selected from Datapro 70.* $12.

☐ **All About EDP Media and Supplies**

This 30-page report covers a wide range of products such as disk packs, magnetic tape, punched cards, computer forms, etc. A total of 164 vendors and their product lines are detailed. Also included are tabulations of the responses from 1870 users who rated disk packs, diskettes, magnetic tape, punched cards, printer forms, and printer ribbons. *Selected from Datapro 70.* $12.

MORE There are more time- and money-saving reports on the next page

How to order:

Simply check the boxes in front of the reports you want to order, fold this complete postage-paid form so that the green Datapro address faces out—staple, and mail. Or, save $2.00 handling charge by sending your check with order. Datapro will send the reports by return mail.

NOTE: **Be sure that your name and address on the lower panel of the back page is complete.**

Please enter my order for _____ reports at $12 each.

☐ Enclosed is $_____.

☐ Bill me for $_____, adding $2.00 handling charge.

☐ Please send me information about Datapro's monthly updated information services.

☐ Please send me information on Datapro's EDP/Office System seminars.

My telephone number is _____.

Form 2413-462-0

datapro A McGraw-Hill Company

Datapro 30-Day Trial Offer

Use any of these cost-saving information services in your office for 30 days.

DATAPRO REPORTS ON
BANKING AUTOMATION

Datapro Reports on Banking Automation is designed to provide authoritative information on the full spectrum of automated banking systems. Objective time- and money-saving evaluations of equipment and software offerings from major vendors, EFTS equipment, automated tellers, MICR equipment, teller terminals, credit authorization and application software make it easy to compare product features, prices, and performance. Also includes case histories, users' ratings, "How to" articles, and management summaries. A must for today's progressive banker.

Annual subscription price: $290.

DATAPRO REPORTS ON
MINI-COMPUTERS

Datapro Reports on Minicomputers covers all aspects of the fast growing mini/micro computer industry. Contains more than 2,000 pages of detailed product descriptions, specifications, case histories, users' ratings, and objective evaluations of microprocessors, microcomputers, minicomputers, small accounting computers, software, services—and the companies that provide them. Makes it easy to compare product cost/performance. A vital tool for every segment of the data processing industry including systems designers, end users, and equipment vendors.

Annual subscription price: $430.

DATAPRO REPORTS ON
OFFICE SYSTEMS

Datapro Reports on Office Systems covers the full spectrum of office products, systems, and techniques, including: word processing systems, dictation equipment, copiers, data processing services, microform systems, telephone and voice communications systems, addressing and labeling machines, facsimile devices, and calculators. Contains hundreds of product/price comparison charts, users' ratings, management summaries, "How to" articles, and case histories. This three-volume reference is a vital investment in time and cost savings.

Annual subscription price: $430.

THE EDP BUYER'S BIBLE
DATAPRO 70

Datapro 70 is the world's most widely used EDP information service with well over 10,000 users worldwide. This three-volume continually updated general reference service provides more than 2,300 pages of case histories, users' ratings, management summaries, and independent evaluations of EDP hardware, software, services, and suppliers. Hundreds of easy-to-read charts comparing product features, specifications, prices, and performance make Datapro 70 essential for anyone involved with EDP, including buyers, specifiers, planners, designers, and vendors.

Annual subscription price: $540.

DATAPRO REPORTS ON
RETAIL AUTOMATION

Datapro Reports on Retail Automation is the most comprehensive work ever published on current POS/retail automation equipment and systems. Offers more than 400 pages of detailed reports and comparison tables on integrated POS systems, electronic cash registers, EFTS, credit and payment systems, vendors, applications, specialized equipment, and software. An important time- and cost-saving service for buyers, product planners, specifiers, marketers, merchandisers, system designers, and others.

Annual subscription price: $290.

DATAPRO DIRECTORY OF
SOFTWARE

Datapro Directory of Software is a new and better way to review and compare the industry's available software products for a wide range of applications. Contains thousands of objective, uniformly written software descriptions, users' ratings, "make or buy" criteria to put price/performance ratings in perspective, product history, number of users, time-sharing availability listings, and hardware/system requirements. Datapro Directory of Software is your best source of reliable, cost-saving software information.

Annual subscription price: $270.

DATAPRO REPORTS ON
DATA COMMUNICATIONS

Datapro Reports on Data Communications is a totally different service offering broader scope and greater depth on the complete universe of data communications products, services and techniques. This three-volume service contains product profiles, comparison charts, and users' ratings on communications processors, software and terminals, and much more. Plus management guidelines, concepts, and tutorials. Important for the data comm user and vendor.

Annual subscription price: $390.

DATAPRO
EDP SOLUTIONS

Datapro EDP Solutions is an invaluable reference service that provides authoritative and well-documented answers to challenges facing EDP management. The two-binder service offers tried and true solutions to unavoidable problems of personnel management, technology awareness, systems development, operations optimization, standards, and procurement. Not technical in its approach, Datapro EDP Solutions delivers time- and money-saving answers to the EDP manager responsible for planning, designing, programming, or directing a computer installation.

Annual subscription price: $190.

Annual Subscription includes: 1: Complete looseleaf volumes 2: Twelve monthly report supplements 3: Twelve monthly newsletters 4: Telephone inquiry service.

For $10, you can try any one of these 8 Datapro services for 30 days. With no purchase obligation.

Nothing we can say—not even this informative report—can convey the full nature, scope and value of Datapro's total information services. However, by taking advantage of this special $10 offer, you can *experience firsthand* the time- and money-saving benefits of these services.

For just $10, we will put any of these looseleaf services in your hands for 30 days, and give you use of our "Hotline" Telephone inquiry Service. As a special Trial Subscriber, you will enjoy all of the privileges of a regular Datapro Subscriber—at a fraction of the cost. You are under no obligation beyond the $10 review fee and a promise to return the books if you decide not to continue the subscription.

Write or call Datapro today to initiate your 30-day trial subscription. See for yourself—as thousands of others do each month—how the value and information you get with Datapro information services will repay your $10 trial investment many, many times.

For those who value information.

datapro

Datapro Research Corporation □ 1805 Underwood Boulevard □ Delran, NJ 08075 □ 609/764/0100 □ A McGraw-Hill Company

McGraw-Hill